
AD-A239 804

RL-TR-91-125
Final Technical Report
June 1991



SMART FILTER DESIGN

University of Dayton

David L. Flannery and Steven C. Gustafson



APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

Rome Laboratory
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700



This report has been reviewed by the Rome Laboratory Public Affairs Division (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-91-125 has been reviewed and is approved for publication.

APPROVED:

JOSEPH L. HORNER Project Engineer

FOR THE COMMANDER: Harold Roth

HAROLD ROTH, Director

Electromagnetics & Reliability Directorate

If your address has changed or if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify Rome Laboratory (EROP) Hanscom AFB MA 01731-5000. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

gathering and mansaining the data needed, and cor collection of information, including suggestions for in Davis Highway, Subs 1204, Arlington, VA 22202-4302	educing this burden, to Washington Headqua	rters Services, Directorate fo	rding this burden segimete or eny other aspect of this r information Operations and Reports, 1215 Jefferson plact (0704-0188), Washington, DC 20503.		
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE		PORT TYPE AND DATES COVERED		
•	June 1991	Fi	inal Sep 87 - Apr 90		
4. TITLE AND SUBTITLE			5. FUNDING NUMBERS		
SMART FILTER DESIGN		İ	C - F19628-C-0073		
SPERT FILLER DESIGN			PE - 61102F PR - 2305		
6. AUTHOR(S)	·_····································		PR - 2305 TA - J7		
• •	ouen C Custofeen		WU - 31		
David L. Flannery and Sto	even C. Gustarson		- 32		
7. PERFORMING ORGANIZATION NAM	IE(S) AND ADDRESS(ES)		8. PERFORMING ORGANIZATION		
University of Dayton			REPORT NUMBER		
Research Institute					
Dayton CH 45469			UDR-TR-91-03		
9. SPONSORING/MONITORING AGEN	CY NAME(S) AND ADDRESS(ES)		10. SPONSORING/MONITORING		
Dane Value (Profil			AGENCY REPORT NUMBER		
Rome Laboratory (EROP)	0.1		RL-TR-91-125		
Hanscom AFB MA 01731-500	U				
11. SUPPLEMENTARY NOTES					
Rome Laboratory Project	Engineer: Joseph L. H	orner/EROP/(6	17) 377-3841		
12a. DISTRIBUTION/AVAILABILITY STA	TEMENT		12b. DISTRIBUTION CODE		
Approved for public rele	ase; distribution unli	mited.			
which encode modulation magneto-optical spatial named Metric Sort was concorrelator. In an added approaches for on-line apromising results. The referred journals.	plementation of distort acentrated on ternary produces of -1, 0, and distort. A proceed and successful distort, initial simulation/synthesis of	chase-amplitude, and their is cowerful new so ly demonstrate lations invest correlation published pap	e filters (TPAF), mplementation using mart TPAF formulation ed on a laboratory igating neural network filters yielded ers including seven in		
14. SUBJECT TERMS Optical Correlation, Pa	attern Recognition. Au	gtomatic Targe	15 NUMBER OF PAGES 124		
Recognition		-	16 PRICE CODE		
	18. SECURITY CLASSIFICATION	19. SECURITY CLAS	SIFICATION 20. LIMITATION OF ABSTRACT		
OF REPORT	OF THIS PAGE	OF ABSTRACT UNCLASSIF			
UNCLASSIFIED	UNCLASSIFIED	UNOLROOT!	Standard Form 298 (Rev. 2-8)		



Acces	sion Fo	F	l
NTIS	GRA&I		
DTIC	TAB	Ď	
Unann	ounced		
Justi	ficatio	n	_
			_
Ву			
Distr	ibution	1/	
Availability Codes			
	Avail	and/or	
Dist	Spec	ial	
1	!!!		
n'			
•			

Contents

1	Inti	roduction	1
2	Smart Ternary Phase-Amplitude Filters		
	2.1	Initial TPAF Concept and Simulations	3
	2.2	Basic Design Elements of BPOFs	4
	2.3	Initial Smart TPAF Concept - the fTPAF	4
	2.4	Experimental Implementation of Ternary Filter Modulation	5
	2.5	Metric Sort Smart TPAF Formulation - analytical and experimental results	6
3	Neu	ral Networks for Synthesis or Adaptation of TPAFs	11
	3.1	Work Addressing Adaptation of Binary Amplitude Masks	12
	3.2	Distortion Parameter Estimation	12
	3.3	Binary-phase Filter Synthesis	13
4 Conclusions and Recommendation		iclusions and Recommendations	14
	4.1	Smart Filter Design	14
	4.2	Neural Networks for Filter Synthesis	15

List of Figures

2.1	Diagram of	128x128 elem	ent MOSLM	correlator system	1	10
-----	------------	--------------	-----------	-------------------	---	----

Section 1

Introduction

The is the final technical report for a basic research effort entitled "Smart Filter Design" performed by the Applied Physics Division of the University of Dayton Research Institute (UDRI) during the period from September 1987 to April 1990. The effort was sponsored by the U.S. Air Force Rome Air Development Center, Cambridge Research Laboratories (RADC/ESO). Dr. Joseph Horner was the Air Force program manager.

The central theme of the effort was the investigation of the ternary phase-amplitude filter (TPAF) correlation filter concept, which was proposed as a practically implementable improvement of the binary phase-only filter (BPOF) that had superior discrimination properties. The effort included both experimental work to demonstrate practical implementation of TPAFs using existing spatial light modulators (SLMs) and analysis and simulation to investigate the performance and design of smart TPAFs. Excellent results were obtained in all areas of investigation including the following:

- Analytical and matching experimental demonstrations of improved correlation discrimination using TPAF formulations.
- Experimental implementation of TPAFs using magneto-optic SLMs.
- Development of two effective TPAF smart filter formulations.
- Experimental demonstration of "metric sort" smart filters applied to a challenging target recognition problem with results in good agreement with simulation predictions.

The Smart Filter Design effort was augmented in May 1989 to include the investigation of neural network approaches to correlation filter adaptation and synthesis in added effort sponsored both by RADC/ESO and the U.S. Army Missile Command, Redstone Arsenal, Alabama. Dr. Donald Gregory was the Army program manager. This portion of the effort has generated favorable analytical results which demonstrate the potential of neural networks to adapt or synthesize TPAF filters that optimize correlation in response to varying input scene clutter and target distortions.

The Smart Filter Design contract effort has produced results reported in 16 publications, including 7 articles in major refereed journals and 9 in conference proceedings [1]-[16]. These papers serve to document the effort and they will be used to advantage in this report. Many of them are included as appendices.

This report is organized as follows. Section 2 presents TPAF correlation filter work and section 3 discusses neural network-related efforts. Section 4 provides conclusions and recommendations for future investigations.

Section 2

Smart Ternary Phase-Amplitude Filters

This section traces the evolution of TPAF formulations from the basic concept originally proposed, which addressed only improved discrimination and had not been experimentally verified, to the culmination of the contract effort, when metric sort smart TPAFs were experimentally and successfully applied to a complex and realistic target recognition problem. The discussion of each research advance will rely heavily on published papers included as appendices to this report.

2.1 Initial TPAF Concept and Simulations

The TPAF was conceived and proposed for the Smart Filter Design program as a significant improvement of the BPOF, that involved the addition of zero modulation (i.e., an opaque or blocking element) resulting in three discrete filter modulation levels: -1, 0, and 1. The TPAF was expected to provide enhanced discrimination against noise, clutter, and other nontargets in the input scene by selectively blocking spatial frequencies containing mostly nontarget information. A key feature motivating the TPAF concept was the prospect of easy near-term implementation using available SLMs, as was already demonstrated in the case of BPOFs using magneto-optic SLMs (MOSLMs) [2,17]. Our proposal speculated that the MOSLM devices also could implement TPAF modulation levels. An additional intrinsic benefit due to the discrete nature of the filter is efficient electronic storage of filter patterns, which require at most two bits (binary digits) of storage per filter element or twice that of a BPOF.

Our proposal provided 64x64 element simulations that indicated improved discrimination of TPAFs for simple binary images representing a gunboat and a "gunless" boat

immersed in random binary noise. These results were subsequently published [1,6] and copies of the papers are included as Appendices A and B. This work introduced the transform-ratio technique for setting the zero-state pattern (called the binary amplitude mask or BAM) of a TPAF. The phase states were set as in a BPOF by thresholding on the phase of the Fourier transform of the reference pattern.

Thus, at the start of our effort the TPAF had been defined and shown in simulations using the transform-ratio BAM design technique to provide improved discrimination over a BPOF. Experimental TPAF implementations and smart TPAF formulations remained to be addressed.

2.2 Basic Design Elements of BPOFs

A study of the basic design elements of BPOFs, threshold line angle (TLA) and reference offset, was believed necessary to support proper design of either BPOFs or TPAFs. The study was conducted by computer simulation of correlations using 256x256 element arrays containing scenes based on actual visible tank and background terrain images. Significant but relatively small variations of correlation performance were found as the design parameters were varied, and the results did not support the choice of a unique optimum set of parameters. Rather it was judged that the choice of design parameters must be considered case specific and thus explored for each application.

Related analytical work showed that the TLA, while varying continuously over the nonambiguous range of 0 to 90 degrees, defined, at specific angles, the previously reported BPOF types, including cosine, sine, and Hartley BPOFs. Analysis also defined an equivalent POF (phase-only filter) that can be synthesized for any BPOF made with a particular TLA. This concept can aid in predicting the impulse response of a BPOF.

Our work on BPOF design elements was reported in October 1988 [7], and a reprint of that paper is included as Appendix C.

2.3 Initial Smart TPAF Concept - the fTPAF

The first smart TPAF concept investigated was an extension of a BPOF smart filter concept reported by Jared and Ennis [18,19]. Their BPOF was based on thresholding the transform of a composite (training) image set with iterative adjustment of superposition weights to achieve distortion invariance across the set. The ad hoc numerical procedure

may not converge but does so in many cases of interest. Our modification was to impose a BAM on the BPOF, thus creating a TPAF. The BAM was defined by the transform-ratio technique applied to representative composite power spectra of the in-class and out-of-class training sets. It was held fixed as the weights defining the binary-phase states were iterated. We named our variant the fTPAF in keeping with the nomenclature used by Jared and Ennis, e.g., the fSDF-BPOF.

The fTPAF concept was first reported in April 1988 [3] along with initial simulations indicating promising performance for simulated images of two types of wrenches. A reprint of this paper is included as Appendix D.

A much more extensive set of simulations applied the fTPAF to a complex character recognition problem. This work was reported in May 1989 [9], and a reprint of the paper is included as Appendix E. The iterative adjustment of weights to form a filter involves evaluating correlations of the current composite TPAF with all in-class training images at each step. Jared and Ennis simplified this phase of the computation by examining only the central (on-axis) output location where the correlation peak is expected to appear. This enables a simpler inner product operation to be used instead of the inverse Fourier transform required to generate a complete correlation pattern. In our character recognition work we performed the full correlation and based our weight iterations on the highest peak in the output regardless of location. We believe this procedure is necessary for the character recognition problem because the response anywhere within a character block is used as the basis of discrimination, and cross-correlations of characters are expected to have considerable side-lobe structure.

The reported work established the fTPAF as an effective smart TPAF formulation having good overall performance and providing controllable tradeoffs of distortion invariance and nontarget discrimination.

2.4 Experimental Implementation of Ternary Filter Modulation

Results published by workers at Litton Data Systems [20] indicated that a third MOSLM modulation state, called the "mixed magnetization" or neutral state, was easily accessible. We conducted experiments with a Semetex 128x128 MOSLM in a coherent imaging setup to investigate the modulation characteristics of this state. We found that reasonably effective zero-state modulation was easily achieved with slight modifications of the software driving the MOSLM addressing circuitry. The intensity transmission of the

"off" state using a Nyquist optical system bandpass (based on the SLM pixel spacing) was measured to be no more than 3% of the "on" (i.e., ± 1 modulation) pixels. The Nyquist bandpass of the filter SLM corresponds to the output plane extent in a typical correlator design. These results indicated good potential for practical TPAF implementation, although the 3% leakage is not ideal. Our findings were first announced at the Annual Meeting of the Optical Society of America in November 1988 [21], and were then published in a letter in Applied Optics in March 1989 [8]. The letter was jointly authored with researchers from New Mexico State University, who reported experiments demonstrating improved correlation SNR using a low-frequency block implemented with the third state in a MOSLM filter. A copy of this letter is included as Appendix F.

A series of experiments were conducted to investigate implementation of transformratio TPAFs with significant simulation-predicted discrimination improvements. These filters involved BAMs (i.e., zero-state patterns) having much more complicated structures than the simple low-frequency block previously implemented by the workers at New Mexico State University (just discussed). The work was conducted on our 48x48 element MOSLM correlator [2,17] and involved discrimination between various binary targets and nontargets. (Distortion invariance was not addressed in these experiments.) Generally excellent agreement between simulation predictions and experimental results was observed in cases involving significant improvements in discrimination between shapes (e.g., from two dB for a BPOF to five dB with a TPAF). An extensive simulation/experimental investigation addressed variations of the percentage of filter elements blocked, TLA variations, and the effect of zero-state leakage polarity relative to phase-state polarity. It was found that zero-state leakage (about 3% in intensity as mentioned) limited useful TPAF performance to filters having less than 70% of their elements set to zero. This work was first reported at an SPIE conference in August 1989 [10], and appeared in Optical Engineering in May 1990 [15]. Copies of these papers are provided as Appendices G and H.

2.5 Metric Sort Smart TPAF Formulation - analytical and experimental results

Although the fTPAF is an effective smart TPAF formulation, it has been superceded by a new formulation called "metric sort" (or MS for brevity) developed early in 1990 on the Smart Filter Design program. The MS formulation is superior overall based on simulation results to be shown and also requires less computational effort.

The MS filter formulation uses an iterative numerical process which attempts to maximize the classical SNR of a filter constrained to TPAF modulation values. A numerical

procedure for solving this problem for the case of additive Gaussian white noise has been given recently by Kumar and Bahri [23]. Our treatment differs from theirs in at least two important ways: (1) we treat any noise spectrum, removing the need to assume white noise, and (2) whereas their algorithm was exact and exhaustive (i.e., guaranteed to find the global maximum SNR) ours involves approximations and is subject to trapping to local SNR maxima depending on an assumed starting point. The above comments apply to the case of a single target pattern with additive stochastic noise. To use the metric sort algorithm as a smart filter, we make the ad hoc extensions of (1) using a composite in-class target image to impose distortion invariance over a training set and (2) applying the filter to scenes in which noise is not additive (because background noise does not invade the target). Nevertheless we have obtained better results for a challenging correlation problem to be discussed than with several other smart BPOF and TPAF formulations.

The classical SNR metric addressed by the MS procedure is:

$$SNR = \frac{\mid \int S(f)H(f)df\mid^2}{\int N(f)\mid H(f)\mid^2 df}, \qquad (2.1)$$

where f is the (spatial) frequency variable and S, H, and N are the signal Fourier transform, filter function, and noise spectrum, respectively. A brief description of the MS formulation algorithm follows:

- 1. Assume a starting TPAF pattern (e.g., a BPOF formed with a particular TLA) and calculate and store the corresponding SNR (Eq. 2.1).
- 2. For each filter pixel, calculate the differential SNR resulting from switching the pixel to each of the remaining two allowed filter modulation states (-1, 0, 1). Store the maximum of these two values and the associated new modulation state.
 - 3. Sort the pixels into descending order of the SNR differentials obtained in step 2.
- 4. For the top M values on the sorted list, make the associated pixel modulation state changes and update the SNR value. If a negative SNR differential is encountered, stop and go to step 6.
- 5. Go to step 2, except if an operator-selected maximum number of pixels to be blocked has been exceeded, go to next step.
 - 6. Write the new TPAF filter to a file and end.

Several comments apply to this procedure. First, it clearly has the general nature of a gradient-descent or hill-climbing algorithm. Thus it should be subject to trapping at local

maxima and sensitivity to starting point assumptions. Second, an ideal implementation would require M=1, i.e., resorting the differentials after each single pixel is switched to a new modulation state. This would be extremely compute-intensive, and we suspect not necessary, provided M is very small compared to the number of filter elements, N. The fractional change in SNR in response to changing a single element (out of, for example, 16,384 in a 128-by-128 element filter) is very small. Typically we have used 0.05N as the maximum value for M. The reason for the exception in step 5 is that empirical evidence indicates that filters which optimize SNR are frequently deficient in other metrics such as peak height (Horner efficiency) or peak width. Wider peaks (i.e., poorer location estimates) and lower efficiency typically are associated with increased fractions of pixels blocked. Thus we allow for the possibility of stopping the iterations before convergence to preserve better overall performance, based on the filter designer's judgment.

We now discuss an example of the application of the MS filter to a complex, demanding, and realistic target-in-background problem involving laser radar (LADAR) range images from actual sensors. The noise spectrum used in MS filter formulation was a composite of power spectra from representative background scenes. The signal transform was the Fourier transform of a weighted composite of training set images comprising rotations of the target object. Uniform composite in-class weights were used.

This application of the MS filter is detailed in Appendix I, which is a copy of a paper presented in April 1990 [14]. A MS filter was designed to recognize a target helicopter in the presence of background noise and two other (nontarget) vehicles over a rotation range of ±10 degrees. A simple BPOF provided 6.2 dB SNR for an aligned target, and performance degraded by 3 dB for ±1.5 degrees rotation of the input target. The computer-simulated SNR (called peak-to-clutter in Appendix I) and peak intensity (in essence, Horner efficiency) performance of the MS filter are plotted in Figures 6 and 7 of the article, along with the same metrics for the other smart BPOF and TPAF filters studied for comparison. The MS filter is clearly superior to the other types and also meets the design goals (good performance over the design rotation range). Note in particular the inferior performance of the fTPAF which was the best smart TPAF formulation known prior to the MS filter. Note also the even poorer performance of a composite BPOF. Although this filter did not benefit from the use of techniques such as the successive-forcing-algorithm [22] or Jared-and-Ennis (fSDF-BPOF) formulation to adjust weights or allow for complex weights, it seems unlikely that performance could equal the MS filter even with these improvements. The need for blocking elements in a filter is dramatically illustrated by the poor SNR performance of this filter. (See Figure 6 of Appendix I.)

In the work reported in Appendix I, five different initial TPAFs were tried as starting points for the MS formulation, comprising BPOFs formulated with TLAs of 0, 22.5, 45,

67.5, and 90 degrees. Best overall performance was obtained by stopping iterations when about 54% of filter elements had been switched to the zero-state. The five resulting filters showed slight variations but approximately equal overall performance for all TLAs except 90 degrees, which was noticeably inferior. The filter resulting from a starting TLA of 45 degrees was chosen as representative of best overall performance and was used in the work reported.

In one case, TLA = 0, the MS filter was iterated to convergence, resulting in 92.5% of the elements being set to the zero-state. This filter exhibited only slightly better P/C than the tested MS filter, although uniformity over the ± 10 degree testing interval was significantly improved. These improvements were obtained at the expense of a 4.3 dB reduction in Horner efficiency, to 0.29%, or a 9.8 dB reduction from the efficiency of the BPOF. The practical impact of this depends on the specific correlator hardware implementation. It probably still compares favorably to the efficiency of a classical matched filter.

In Appendix I, Figure 8 and the preceding table provide experimental results in good agreement with simulations. These results were obtained on a laboratory correlator using 128x128 Semetex MOSLMs. The correlator was assembled and first operated on the Smart Filter Design program. Figure 2.1 is a diagram of the experimental setup. The major components of this correlator were furnished by Martin Marietta Corporation, Strategic Systems, Denver, Colorado, with the exception of the input SLM, which was furnished by UDRI as cost sharing on the Smart Filter Design contract.

Although the MS filter (currently) optimizes only with regard to a classical SNR definition, which is not necessarily best for many practical correlation problems, its excellent performance is believed to result from three key factors:

- It can handle any noise spectrum (not just white noise).
- It directly optimizes each filter element (rather than parametrizing to a dimensionality much less than the filter space-bandwidth product).
- It implicitly generates TPAF modulation levels (rather than quantizing a continuous-valued function).




Figure 2.1: Diagram of 128x128 element MOSLM correlator system

Section 3

Neural Networks for Synthesis or Adaptation of TPAFs

This work is motivated by a long-term high-payoff vision: by synthesizing correlation filters, a trained neural network could assume all or most filter storage and selection functions for an adaptive optical correlator. This vision is reasonable on basic physical grounds because, although the Fourier transform operations performed by an optical correlator are appropriate for translation processing (e.g., for locating a target), it is well known that without suitable coordinate transformations these operations cannot be directly applied to distortion processing (e.g., for determining target rotation or scale). Thus it is logical to envision a system in which an optical correlator determines target location and in which an associated executive processor determines target rotation, scale, aspect angle, and other distortions and uses this distortion information to synthesize updated or adapted correlation filters and to control correlator operation.

Supervised-learning neural networks are particularly appropriate for this executive processor because they can learn complex tasks (such as target aspect angle determination) and, once trained, can perform their learned processing tasks in real time. The result, particularly if hardware (and eventually optical hardware) instead of software-simulated neural networks is used, could be a compact and efficient system that in synthesizing correlation filters has, in effect, the equivalent of a large correlation filter storage capacity with automatic interpolation and extrapolation (due to associative-memory-accessible, fault tolerant, and parallel and distributed patterns of network interconnection strengths). This neural network system could control the optical correlator by synthesizing the filters used to determine target location so as to compensate or adapt for evolving target distortions.

In the initial work performed on the Smart Filter Design effort and reported here, TPAF formulation was viewed as two separate processes: binary-phase state formulation (with the primary objective of adapting to evolving target distortions) and binary amplitude mask (BAM) formulation (with the primary objective of optimizing SNR for changing input scene noise).

3.1 Work Addressing Adaptation of Binary Amplitude Masks

Our first foray into filter synthesis involved using a neural net to set the crude bandpass characteristics of a BPOF (binary phase-only filter), which thus comprised a TPAF (ternary phase-amplitude filter) by definition. This work is described in detail in a paper presented in October 1989 [12] and included here as Appendix J. Four bandpass rings spanning most of the Fourier plane were used, constituting a simple example of the binary amplitude mask (BAM) component of a TPAF. The network inputs were four power spectrum features obtained by applying the same bandpass rings to the spectrum of the input scene. The outputs were four binary values which established the on/off status of the respective BAM rings. Computer simulations were performed using 64x64 element arrays with binary input scenes consisting of an airplane shape in random noise backgrounds of varying energy levels. The results were promising: the backpropagation-trained network was able to adaptively set the filter bandpass rings so as to achieve signal-to-noise performance comparable to that of the best possible BAM-ring combination (determined by exhaustive trials) as the input noise level was varied from zero to a maximum level.

This initial effort verified in part an intuitive concept, which is that input scene power spectrum information is a good basis for determining the binary amplitude states of a correlation filter. It also demonstrated that a neural net can formulate near-optimum (albeit elementary) BAM patterns for simple test cases.

3.2 Distortion Parameter Estimation

Neural network techniques were investigated for estimating target distortions (e.g., rotation and/or scaling) based on the shape or sidelobe structure of BPOF/TPAF correlation peaks.

The inputs to the backpropagation-trained neural network were taken from a 5x5 sampling grid centered on the target peak resulting from correlation with a known reference filter. The network output was an analog value representing an estimate of a distortion parameter (e.g., rotation angle).

Initial simulations using 128x128 element arrays and binary targets with no background or noise (see Figure 1 of Appendix K) showed very good results, which were reported in April 1990 [16]. A copy of this paper is included as Appendix K. BPOFs with quarter-bandpass BAMs (amplitudes set to zero for frequencies outside a bandpass equal to one quarter of the Nyquist bandwidth) were used in these simulations. Figures 2-4 of Appendix K are plots of rotation angle estimation accuracies from these simulations.

3.3 Binary-phase Filter Synthesis

Filter phase-portion synthesis was investigated using an approach based on the same correlation-peak shape (or sidelobe) sampling technique used for distortion parameter estimation (see above), except that the network output was a considerable number of binary values which set the phase states of a BPOF.

The use of reduced-format filters is desirable to keep the number of neural net outputs at a practical level in view of available computer simulation resources. Our first efforts used polar formatted filters having 32 sectors, of which only 16 were independent since the filter was trained as a cosine-BPOF. The results were promising [16] (see Figures 11 and 12 in Appendix K). However, performance of the synthesized filters was limited by the small number of available free parameters. Filters were synthesized over a 45-degree range and performed as well as the training set filters with no background noise or other nontarget shapes in the input scene. Correlation performance degraded significantly for even small amounts of background noise added to the input scenes.

Section 4

Conclusions and Recommendations

4.1 Smart Filter Design

The work performed on the Smart Filter Design effort and reported here constitutes the successful development of the TPAF concept from a rudimentary idea to a full fledged smart filter concept with experimentally demonstrated effectiveness in addressing realistically complex automatic target recognition goals. The TPAF provides significant advantages over the BPOF (and we strongly believe over any phase-only filter), particularly with regard to nontarget discrimination. The TPAF advantages are obtained with little additional expense or complication. The TPAF can be implemented with an available modulator at attractive frame rates, and it enables efficient electronic storage of filter patterns. The Metric Sort TPAF implemented with available magneto-optic SLMs comprises a technology base for near-term correlation systems to address automatic target recognition and other military pattern recognition problems.

Several recommendations can be made to extend and perfect the current state of the art:

- The Metric Sort formulation should be modified to address other metrics, including control of peak width and shape or sidelobes and rejection of specific nontarget shapes.
- An improved global optimization technique (e.g., simulated thermal annealing) should be investigated as a modification of the Metric Sort algorithm.
- In anticipation of the practical availability of phase-only electrically addressed SLMs which could be used with cascaded binary amplitude modulators (or might

incorporate zero-state modulation directly), modifications of TPAF theory to incorporate (and presumably realize advantages from) continuous phase modulation should be investigated. This filter type might be called the "continuous-phase binary-amplitude filter" (CPBAF).

- The continued development of magneto-optic SLM technology is strongly recommended. The issue of good zero-state performance (to support TPAF implementation) should be addressed.
- The development of other SLM types to support TPAF modulation or to provide efficient binary-amplitude modulation in tandem with a continuous phase-only modulating SLM should be considered or continued as appropriate.
- In connection with the previous recommendation, the effects of using reduced spatial resolution in the binary-amplitude mask of a TPAF (relative to the phase portion) should be investigated. It seems possible that resolution of the BAM can be lower than the BPOF while maintaining most of the advantages of ternary modulation in some cases.

4.2 Neural Networks for Filter Synthesis

The work performed on this effort has demonstrated potential for neural networks to formulate or adapt TPAFs in response to varying input scene noise and target distortions. The reported work was limited to test images of little practical importance, but the principles investigated should be applicable to more challenging cases that will require primarily more extensive neural networks.

Recommendations center on extending the current work to address more realistic and challenging inputs and filter geometries as follows:

- Filter patterns determined by larger numbers of free parameters should be investigated (e.g., 500). These may provide better correlation performance for realistic scenes.
- Additional input signals for the neural network should be investigated, including
 input-scene power-spectrum features obtained by a sampling pattern more complex
 than the four bandpass rings used in the reported work. Alternative schemes for
 sampling in the neighborhood of correlation peaks should be considered. Sampling
 from the input scene itself or from characteristics such as the intensity histogram
 should be considered.

• The correlation peak sampling technique used to provide neural net inputs for filter phase-portion synthesis appears to be fundamentally limited by the requirement to correlate with a predetermined reference filter (upon which the network was trained) to generate the required correlation peak. This leads to a "bootstrapping" problem: the new synthesized filter is not (in any practical sense) predetermined and thus cannot be used to train a network. Thus the filter synthesis process breaks down after one cycle of correlate-sample-synthesize, i.e., the system cannot follow the target through significant distortions (those sufficient to preclude good signal-to-noise performance of the original reference filter). Methods to circumvent this limitation should be explored.

Bibliography

- [1] D. Flannery and S. Cartwright, "Optical Adaptive Correlator," Proc. Third Annual Aerospace Applications of Artificial Intelligence Conference, pp. 143, Dayton, Ohio, October, 1987.
- [2] D. Flannery, J. Loomis, M. Milkovich, and P. Keller, "Application of binary phase-only correlation to machine vision," Optical Engineering, Vol. 27, pp. 309-320, April, 1988.
- [3] D. Flannery, J. Loomis, and M. Milkovich, "New formulations for discrete-valued correlation filters," Proc. SPIE, Vol. 938, pp. 206-211, April, 1988.
- [4] D. Flannery, J. Loomis, and M. Milkovich, "Application of Binary Phase-only Filters to Machine Vision," Proc. SPIE, Vol. 960, pp. 114-130, June, 1988.
- [5] D. Flannery, J. Loomis, and M. Milkovich, "The Use of Binary Magneto-optic Spatial Light Modulators in Pattern Recognition Processors," Optical Society of America Technical Digest Series 1988, Vol. 8, p. 221, June 1988.
- [6] D. Flannery, J. Loomis, and M. Milkovich, "Transform-ratio ternary phase-amplitude filter formulation for improved correlation discrimination," Applied Optics, Vol. 27, pp. 4079-4083, 1 October 1988.
- [7] D. Flannery, J. Loomis, and M. Milkovich, "Design elements of binary phase-only correlation filters," Applied Optics, Vol. 27, pp. 4231-4235, 15 October 1988.
- [8] B. Kast, M. Giles, S. Lindell, and D. Flannery, "Implementation of Ternary Phase-amplitude Filters Using a Magneto-optic Spatial Light Modulator," Applied Optics, Vol. 28, pp. 1044-1046, 15 March 1989.

- [9] M. Milkovich, D. Flannery, and J. Loomis, "A Study of Transform-ratio Ternary Phase-Amplitude Filter Formulations for Character Recognition," Optical Engineering, Vol. 28, pp. 487-493, May 1989.
- [10] S. Lindell, D. Flannery, "Ternary phase-amplitude filters for character recognition," Proc. SPIE, Vol. 1151, pp. 174-182, August 1989.
- [11] D. L. Flannery and J. L. Horner, "Fourier Optical Signal Processors," Proc. IEEE, Vol. 77, pp. 1511-1527, October 1989.
- [12] D. L. Flannery, S. C. Gustafson, and D. M. Simon, "On-line Adaptation of Optical Correlation Filters," Proc. Fifth Annual Aerospace Applications of Artificial Intelligence Conference, Vol. II, pp. 185-196, Dayton, Ohio, 23-27 October 1989.
- [13] D. Flannery, "Ternary Phase-amplitude Correlation Filters for Real-time Pattern Recognition," Proc. 23rd Asilomar Conference of Signals, Systems and Computers, Vol. 1, pp. 314-317, 30 October 1989.
- [14] D. Flannery, W. Hahn, and E. Washwell, "Application of Ternary Phase-amplitude Correlation Filters to LADAR Range Images," Proc. SPIE, Vol. 1297, pp. 194-206, April, 1990.
- [15] S. Lindell and D. Flannery, "Experimental Investigation of Transform Ratio Ternary Phase-amplitude Filters for Improved Correlation Discrimination," Optical Engineering, Vol. 29, pp. 1044-1051, September 1990.
- [16] S. C. Gustafson, D. L. Flannery, and D. M. Simon, "Neural Networks With Optical Correlation Inputs for Recognizing Rotated Targets," Proc. SPIE, Vol. 1294, pp. 171-179, Orlando, Florida, April 1990.
- [17] D. Flannery, A. Biernacki, J. Loomis, and S. Cartwright, "Real-time Coherent Correlator Using Binary Magnetooptic Spatial Light Modulators at Input and Fourier Planes," Applied Optics, Vol. 25, pp. 466, 1 February 1986.
- [18] D. Jared and D. Ennis, "Inclusion of filter modulation in synthetic-discriminant function filters," Applied Optics, Vol. 28, pp. 232-239, 15 January 1989.
- [19] D. Jared, D. Ennis, and S. Dreskin, "Evaluation of binary phase-only filters for distortion-invariant pattern recognition," Proc. SPIE, Vol. 884, pp. 139-145, January 1988.

- [20] W. Ross, D. Psaltis, and R. Anderson, "Two Dimensional Magneto-optic Spatial Light Modulator for Signal Processing," Optical Engineering, Vol. 24, pp. 485-490, 1983.
- [21] D. Flannery, M. Milkovich, J. Loomis, and S. Lindell, "Optimization of Ternary Correlation Filters for Distortion Invariance," Paper FP5, Optical Society of America 1988 Annual Meeting, Santa Clara, California, 4 November, 1988.
- [22] Z. Bahri and B. V. K. V. Kumar, "Binary Phase-only Synthetic Discriminant Functions (BPOSDFs) Designed Using the Successive Forcing Algorithm," Proc SPIE, Vol. 1297, pp. 188-193, Orlando, FL, 16-20 April 1990.
- [23] B. V. K. V. Kumar and Z. Bahri, "Efficient algorithm for designing a ternary valued filter yielding maximum signal to noise ratio," Applied Optics, Vol. 28, pp. 1919-1925, 15 May 1989.

Optical Adaptive Correlator

David Flannery

The University of Dayton Research Institute Dayton, Ohio 45469

and

Steven Cartwright

The Environmental Research Institute of Michigan Dayton, Ohio 45324

abstract

Real-time analog coherent optical processors generating correlation functions by Fourier techniques have recently demonstrated impressive performance and near term system potential. Correlation is one example of linear shift-invariant filtering which can be performed in two dimensions by such systems. Fourier domain filters are implemented by spatial light modulators, and particularly good results have been obtained with devices which can implement filter functions with complex values limited to a few discrete phase or amplitude levels. More general real-time filtering operations on signals and patterns can be performed on an adaptive basis using hybrid (optical/digital) systems, presumably limited by the generality of filters realizable with optical hardware and the accuracy of the analog coherent optical system. An advanced discrete level filter formulation, derived for correlation purposes, will be reviewed and related to classical optimal filter formulations.

Introduction

Optics was first considered as a computational medium in the 1950's when O'Neill applied linear communications theory to optical propagation [O'Neill, 1956]. Under certain assumptions optics can be viewed as a linear filter, implemented in two dimensions. The advent of the laser provided a bright source of coherent light and gave direct physical access to an object's Fourier transform. The holographic filter allowed complex filtering operations to be performed. More general filters could be implemented through computer-generated holograms. Through the early 1970's many variations and improvements on the holographic matched filter were proposed and implemented,

though few ever saw practical application.

During this period, interest in matched filtering, and optical processing in general, diminished. Such processing ideas were considered interesting but impractical. Electronic computing, on the other hand, continued to make great advances, and was widely embraced for most applications. Flexible nonlinear algorithms and predictable accuracy were fundamental advantages attributed to digital processors, whereas optics was perceived as both inflexible and limited in accuracy by virtue of its analog nature.

In the late 1970's, digital optical computing concepts, in which nonlinear operations were performed on binary data, emerged and received considerable attention. The term "optical computing" was accepted, at least by some, to mean numerical operations on data in discrete representations, as opposed to (analog) optical processing.

The key potential advantage of optical computing is the inherent parallelism of optics, which should lead to massive interconnection capabilities. Promising digital optical computing research efforts are underway, but it is apparent that major success in this area awaits long term developments in several areas, including devices and components for switching and interconnection, and algorithms and control structures suitable for the massively parallel architectures involved. Neural net concepts, a topic of much current interest, may provide an answer to the control and algorithm issue; optics is seen as a likely candidate for implementing the massive interconnections associated with neural nets.

Recent developments have moved the focus of optical computing back towards analog processors. Key device technology improvements have enabled smaller real-time systems. While the resolution of spatial light modulators (SLM) is still not great, new concepts of matched filtering have allowed much better use of the available capacity in filters which encode only two or three discrete levels of phase and/or amplitude modulation.

Coherent optical correlation is one instance of analog Fourier optical processing, in which the Fourier transform of a desired impulse response function is implemented as complex (amplitude and phase) modulation patterns in the Fourier plane (filter plane) of the optical system [Goodman, 1968]. The processor performs general linear shift-invariant convolution in two spatial dimensions, with light amplitude as the analog variable. For correlation the filter function is required to be the complex conjugate of the transform of the desired reference function. The general applicability of this type of optical computing for signal and image processing has been limited by two factors: (1) unavailability of suitable real-time hardware, most notably spatial light modulators, and (2) relatively low accuracy of results, typical of analog processing. In the case of correlation processing, a further limitation has been the inherent shape distortion sensitivity of the correlation process, which may be viewed as a form of algorithmic inflexibility, another negative characteristic frequently attributed to optical processing architectures. This problem may be attacked by faster hardware, and thus is connected with the real-time hardware issue to some extent.

Interest in optical processing, and research and development of related techniques and devices, have continued based on the obvious high performance potential which

can result if the right real-time hardware is applied to an appropriate problem using a well chosen system architecture/algorithm design. The resulting hybrid system, capitalizing on the inherent Fourier transforming power of the optics, would perform at rates, and with power and size metrics, which are superior by orders of magnitude to all-electronic processors.

The hybrid adaptive optical correlator is a leading candidate to achieve an impressive existence proof for useful optical processors as described in the previous paragraph. Research and device development over the past several years, primarily the combination of magneto-optic spatial light modulators [Ross et al., 1982] and the binary phase-only filter (BPOF) concept [Horner and Leger, 1985; Horner and Bartelt, 1985], have led to the laboratory demonstration of real-time correlators [Flannery et al., 1986; Keller et al., 1986] which prove the feasibility of hybrid correlation concepts. Systems which can be constructed in the near term based on these concepts will be capable of performing correlations of 256-by-256 element patterns at rates up to 1000 correlations per second. Each correlation can involve a different input and reference pattern, the latter being stored in digital memory. Such storage is practical and efficient due to the reduced information content of BPOF filters; 1000 such filters will fit in four MB (megabytes) of memory, a rather modest amount with today's digital technology. Rapid access of BPOF patterns is possible because they are stored in the binary form directly useable by the filter SLM; no D/A or other type of conversion is required.

A most notable aspect of the correlation work described is the use of a Fourier domain filter which implements only two possible modulation states: +1 and -1 (or equivalently zero and 180 degrees of phase shift) at each resolvable frequency bin. General filters, including correlation filters, call for continuous independently controlled amplitude and phase modulation, and it was extremely surprising that worthwhile results could be obtained with a two-level phase-only filter function. Only the unavailability of SLMs capable of general complex modulation, combined with the ready availability of relatively well developed binary magneto-optic SLMs can explain the fact that BPOF filters were even considered. Modulators implementing general complex modulation are still unavailable and will be difficult to realize in a single device. Carrier frequency (holographic) encoding currently is the only practical method to implement general complex amplitude modulation, and this technique is wasteful of what little resolution is available in existing SLMs (A factor of five is typical). The input plane device requirement is not restricted by the need for complex modulation since input patterns in most applications are real functions.

The magneto-optic SLMs can implement either binary phase or binary amplitude (on-off) modulation, and are reported capable of ternary modulation with states 1, 0, and -1, equivalent to a cascade of binary amplitude and phase devices. Based on the relatively advanced state of development and commercial availability of these devices, plus the recent successes of BPOF correlation, we are inspired to consider further the application of discrete level filters both for correlation and for general filtering operations. SLMs based on ferro-electric liquid crystals also show development promise and would operate as binary phase or amplitude modulators, with potential for ternary

operation [Clark, 1986; Armitage et al., 1987].

A natural concern is the performance characteristics possible with discrete level filters, since they usually embody only a small fraction of the information content of the full transfer function. This issue has been addressed both theoretically and experimentally for the BPOF used for correlation [Psaltis et al., 1984; Keller et al., 1986], with generally favorable results. We note impressive simulated results obtained by Kallman [1986] with BPOF and ternary filters synthesized for distortion invariant correlation performance, using mini-max optimization techniques. In this paper we summarize correlation results obtained with a new ternary filter formulation, and discuss that filter's relationship to classical optimized filters. This discussion sets the stage for consideration of discrete level filters for more general linear filtering applications.


The Ternary Phase-Amplitude Filter

BPOF filter values are normally assigned by thresholding the reference function (or impulse response) Fourier transform based on a line through the origin in the complex polar plane. One common choice for the angle of this line, for example, corresponds to thresholding on the real part of the transform.

Fourier filters which implement complex modulation states of -1,0, and 1 are labeled ternary phase-amplitude filters (TPAF). We were motivated to consider this extension from the BPOF concept as a potential improvement for correlation use based on a simple ad hoc concept. The BPOF as does any phase-only filter, passes all energy from the input pattern (or signal). This can have advantages, the most notable of which is increased correlation efficiency, i.e. more energy in the correlation peak. However, consider the case in which we wish to construct a "smart" correlation filter; one which will preferentially respond to a desired input pattern (the target) while discriminating against undesired patterns (nontargets including noise).

A simple intuitive concept for improving a BPOF filter is to block (assign a filter value of zero to) those (spatial) frequencies in the filter for which the ratio of target Fourier transform energy to nontarget energy is less than a threshold. The rationale for this is fairly obvious. The spectral areas being blocked are those which tend to pass primarily nontarget energy through to the output (correlation) domain, and thus cannot be contributing favorably to correlation with the desired target. This ad hoc concept is the basis for a smart ternary filter which we label the transform-ratio TPAF.

Figure 1 shows 64-by-64 pixel binary target and nontarget input patterns used in a digital correlation simulation study to assess the promise of the new filter formulation. Random background noise in varying amounts was also added to the input patterns. This noise comprises an additional nontarget pattern, which in this case possessed a flat spatial frequency spectrum. Thus the filter formulation process used two thresholds: one for target-to-nontarget ratio and one for noise blocking. If the target energy in a frequency bin exceeded both thresholds, the corresponding filter element was set to the usual BPOF value; otherwise it was set to zero. The ability of these filters




(b)

Figure 1. Target (a) and Nontarget (b) for Transform Ratio Ternary Smart Filter Study.

to reject noise (S/N) and to discriminate between the desired target and nontarget (S/C or signal-to-clutter) was studied over parametric variations of the two filter design thresholds and amount of input noise. The results are summarized in Figure 2. All filters were formulated using a basic target/nontarget transform magnitude-ratio threshold of 1.1:1, which was known to provide good discrimination performance with zero noise, but was not parametrically optimized.

The trends of Figure 2 are readily apparent. Absolute values are of little significance since they are subject to many arbitrary factors such as resolution and choice of target and nontarget shapes. With little or no noise, filters which use no noise blocking

SIGNAL-TO-NOISE PERFORMANCE OF TPAF FILTER



+0 +0.02 +0.04 filter noise threshold

DISCRIMINATION PERFORMANCE OF TPAF FILTER




Figure 2. Plots of Transform Ratio Smart Filter Performance.

threshold provide good signal-to-noise and discrimination performance. With more noise in the input, the filters with noise blocking thresholds provide better overall performance. Not shown, but an important reference for comparison, is the performance using the simple BPOF formulation based on the target transform. This provided 11.5 dB S/N and 2.57 dB S/C with zero input noise. These values degraded to 1.95 dB and 1.14 dB respectively with 0.4 fractional area noise input. The TPAF formulations with appropriately chosen thresholds are superior in terms of combined noise rejection and nontarget discrimination for any input noise level, and a particular TPAF can be chosen to provide overall superiority over the BPOF over a range of input noise levels. Thus the validity of the transform ratio TPAF concept is supported by these initial simulations. We plan to continue our study of these filters and to implement them experimentally.

Relationship of Discrete-Valued Filters to Classical Optimal Filters

Two classical optimal filters, the Wiener filter and the matched filter, will be compared with the ad hoc BPOF and TPAF filters used in correlation as discussed.

The Wiener filter is designed to provide an optimum (minimum mean square error) estimate of an input signal with additive input noise. It attempts to recover the signal-from the noise. The desired signal and the input noise are assumed to be uncorrelated, and only statistical knowledge of the desired signal is assumed available, i.e., the power spectrum or autocorrelation function. The Wiener filter is defined by:

$$H(s) = \frac{P_{signal}}{P_{nose} + P_{nonel}} ,$$

where s is the frequency variable, P is a power spectrum and H(s) is the filter transfer function. Other assumptions, including ergodicity and that the signal and/or noise have zero mean are implicit in the derivation of this filter function. Wiener transfer functions are real and vary in amplitude from 0 to 1. The filter imposes no phase shifts on any frequency components but simply preferentially weights those areas of the spectrum in which the power spectrum signal-to-noise ratio is higher.

The matched filter is designed to provide optimum detection of an input signal of known form in the presence of uncorrelated additive input noise. The filter produces an output peak with the highest possible ratio to the mean square noise level in the output. It is not designed to recover an estimate of the input signal pattern (waveform or shape), as is the Wiener filter. The form of the matched filter is:

$$H(s) = C \frac{e^{-i2\pi s t_0} S^*(s)}{P_{noise}},$$

where C is an arbitrary complex constant, t_0 is an arbitrary time delay (or shift distance if the filter acts in the space domain), and S(s) is the Fourier transform of the desired

signal. The matched filter can have arbitrary complex gain and an arbitrary time delay or position shift of its impulse response. It imposes phase shifts which are the negatives of those of the desired signal transform. If the noise is white, and the delay and gain constants are taken as unity, the matched filter becomes a constant times the conjugate transform of the desired signal, and thus the filter is performing correlation. Hence optical correlation is often called optical matched filtering.

Although the BPOF and TPAF filter formulations discussed above were synthesized on an ad hoc basis for correlation purposes, and neither is an exact form of a classical optimum filter, it is interesting and instructive to explore their relationship to those filters.

The BPOF filter has unity magnitude, since it is defined as a phase-only filter. It has two discrete phase values which are most commonly obtained by simply thresholding on the phase of the transform of the desired input signal. Thus it encodes desired signal phase quantized to two levels. Clearly it is more akin to the matched filter since that filter incorporates desired signal phase, while the Wiener filter has no phase variation at all. This is not a surprising observation, since the BPOF resulted from an initial goal of correlation, and matched filtering has a close relationship to correlation, as discussed. The BPOF incorporates no magnitude variation, which is a major departure from the classical matched filter formula.

The TPAF filter encompasses the two phase levels of the BPOF and adds a zero level, i.e., it has binary magnitude as well as binary phase values. The TPAF filter is also more akin to the matched filter due to the phase modulation, which simply isn't present in the Wiener filter. The TPAF with its additional degree of freedom (magnitude variation) should enable a better approximation to the true matched filter, which has a definite magnitude variation determined by the ratio of desired signal transform to noise power spectrum. Although derived from ad hoc concepts, the transform ratio TPAF formulation involves a thresholding to determine binary filter magnitude based on the ratio of signal (target) and noise (nontarget) transform magnitudes, and thus would seem to be closely related to the classical matched filter formulation. In addition to the effects of the severe quantization, which are expected to be major, the TPAF filter formulation involves two basic departures from the classical matched filter.

- The matched filter is based on uncorrelated signal and noise. This assumption is violated when the transform ratio technique is applied to highly correlated target and nontarget images such as the boat with and without turret.
- The magnitude of the matched filter is the ratio of the transform magnitude of the desired signal to the noise power spectrum. In the TPAF formulation reported here, thresholding was performed on the ratio of transform magnitudes, which is functionally different. Thus determining the binary TPAF magnitude by thresholding on the quantity

 $\frac{|S|}{P_{\sigma}}$

would correspond more closely to binary quantization of the matched filter magnitude. For the case of white noise as the only nontarget the two ratios give the same results, since the noise power spectrum is then constant.

An additional distinguishing factor is that for binary quantization of phase and magnitude, the design of a TPAF filter involves two adjustable quantization threshold parameters which do not enter into the formulation of the classical matched filter. This is probably fortunate; one can hope that the severe information loss incurred by binary quantization can be compensated to some extent by optimal choice of these two thresholds.

Regardless of whether it is justifiable to view the TPAF as a quantized matched filter, the promising simulation results speak for themselves. Perhaps the proper viewpoint is simply that the transform ratio TPAF is an ad hoc filter formulation which can improve signal-to-noise and discrimination performance over that of the simple BPOF.

We have seen that the discrete filter designs used for coherent optical correlation are (this is not surprising) most closely akin to the classical matched filter. However, a binary filter which is more closely related to the Wiener filter can easily be envisioned and implemented with the same magneto-optic SLMs already used to implement BPOFs, operating in binary amplitude modulation mode. A binary quantized version of the Wiener filter could be implemented with the filter values determined by thresholding on the classical Wiener filter function as determined by estimates of desired signal and noise (nonsignal) power spectra.

Real time adaptive hybrid filtering architectures

Real-time quantized BPOF or TPAF versions of optimal filters can form the basis for powerful hybrid adaptive signal or image processor systems. The generic system concept is illustrated in Figure 3. The figure was prepared with the correlation context in mind, but can be converted to a more general linear filtering context by substituting "rapid sequential convolver" for "rapid sequential correlator".

In the correlation context, the hybrid adaptive system is envisioned attacking the distortion variance nemesis by adaptively selecting from a bank of precomputed smart filter patterns, based on feedback from the correlation output. However, the system can be made more adaptive by incorporating on-line filter synthesis based on short-term estimates of signal and/or noise statistics. As an example, the power spectrum of input patterns is readily available in a coherent optical processor by simply detecting the optical Fourier transform pattern, which is already being generated by the first section of the processor.

Suitable means of gating or selecting input patterns, depending on the particular application, would enable short term averages of P_{ngnal} and/or P_{noise} to be determined and used in filter computations. Where appropriate, the power spectra or transform amplitudes of desired signals or classes of signals and noise would be predetermined

SCHEMATIC OF ADAPTIVE CORRELATOR FOR TARGET RECOGNITION, GUIDANCE, AND MACHINE VISION




Figure 3. Hybrid Adaptive Filtering System Concept.

and stored digitally by the system. Unpredictable or nonstationary signals and/or noise could be handled using the real-time adaptive techniques outlined. The choice of filter type (pseudo-Wiener or pseudo-matched filter) would depend on the system goal (estimation or detection) and the amount of a priori knowledge of the desired signal.

A particularly fortunate match of system capability and problem occurs for the case of short term averaging of power spectra. Averaging over multiple spectra can be accomplished by rapidly sequencing input sample frames while allowing the optically derived power spectrum to time-integrate on a detector surface. The required noncoherent summation is intrinsic to the square law (intensity) response of the photodetection process. A similar process, using movie film as the input, is used in stellar speckle interferometry [Labegrie, 1970]. Current technology trends indicate that the magneto-optic SLMs may be able to achieve frame rates from 3 to 10 times that of typical detector arrays, which meshes quite nicely with the time integration technique.

A means to synthesize discrete value filters having specified impulse responses is of obvious value. Some initial forays in that direction have met with success [Flannery et al., 1987].

Conclusion

We have reviewed the recent advances in optical processing, demonstrated in the context of correlation, which will support practical real-time hybrid adaptive linear signal and image processing systems. The price of practicality, at least in the near term, is the requirement to use filters quantized to binary phase and magnitude levels. Such filters have shown surprisingly good performance for the correlation function, and should have potential for more general convolution processing applications. To view such filters as quantized versions of classical optimal filters is intuitive but not rigorous. The potential for adaptive filtering of images and signals seems great, provided the inherent limitations of discrete level filters and the reduced accuracy of analog optical processing can be dealt with in a particular application.

References

- Armitage, D., J. Thackera, N. Clark, and M. Handschy, "Ferroelectric Liquid Crystal Spatial Light Modulators," Opt. Soc. Am., Proc. Topical Meeting on Optical Computing, Lake Tahoe, NV, p. 221-224 (March 1987).
- 2. Clark, N. and M. Handschy, "Device Application of Ferroelectric Liquid Crystals," SPIE, Optical Computing, Los Angeles, CA, Vol. 625, p. 60-61 (January 1986).
- 3. Flannery, D., A. Biernacki, J. Loomis, and S. Cartwright, "Real-time coherent correlator using binary magneto-optic spatial light modulators at input and Fourier planes," Applied Optics, Vol. 25, No. 4, p. 466 (15 February 1986).
- 4. Flannery, D., P. Keller, S. Cartwright, and J. Loomis, "Characterization of Improved Binary Phase-Only Filters in a Real-Time Coherent Optical Correlation System," SPIE, OE Lase '87, Los Angeles, CA, Vol. 653 (January 1987).
- 5. Gianino, P. and J. Horner, "Additional properties of the phase-only correlation filter," Optical Engineering, Vol. 23, No. 6, p. 695-697 (November/December 1984).
- 6. Goodman, J., Fourier Optics, McGraw-Hill, 1968.
- 7. Horner, J. and H. Bartelt, "Two-bit correlation," Applied Optics, Vol. 24, No. 18, p. 2889-2893 (15 September 1985).

- 8. Horner, J. and J. Leger, "Pattern recognition with binary phase-only filters." Applied Optics, Vol. 24, No. 5, p. 609-611 (1 Mar 1985).
- 9. Kallman, R. R., "Optimal low noise phase-only and binary phase-only optical correlation filters for threshold detectors," Applied Optics, Vol. 25, No. 23, p. 4216-4217 (1 Dec 1986).
- 10. Keller, P., D. Flannery, S. Cartwright, and J. Loomis, "Performance of Binary Phase-Only Correlation on Machine Vision Imagery," SPIE, Optics, Illumination, and Image Sensing for Machine Vision, Cambridge, MA, Vol. 728, p. 257-265 (October 1986).
- 11. Labegrie, A., "Attainment of Diffraction-limited Resolution in Large Telescopes by Fourier Analysing Speckle Patterns in Stellar Images," Astronomy and Astrophysics, Vol. 6, p. 85 (1970).
- 12. O'Neill, E., "Spatial Filtering in Optics," IRE Trans. on Information Theory, Vol. IT-2, p. 56-65 (1956).
- 13. Psaltis, D., E. Paek, and S. Venkatesh, "Optical image correlation with a binary spatial light modulator," Optical Engineering, Vol. 23, No. 6, p. 698-704 (November/December 1984).
- 14. Ross, W., D. Psaltis, and R. Anderson, "Two-dimensional magneto-optic spatial light modulator for signal processing," SPIE Vol. 341, "Real-time Signal Processing V," San Diego, CA, p. 191-198 (August 1982).

Transform-ratio ternary phase-amplitude filter formulation for improved correlation discrimination

David L. Flannery, John S. Loomis, and Mary E. Milkovich

A method of formulating ternary-valued (=1.0.1) correlation filters based on the ratio of spectral energies of target and nontarget patterns is proposed, and performance of such filters is investigated by computer simulations of correlation. The results confirm the intuitive expectation that such filters can enhance signal to-clutter and discrimination performance for target recognition in the presence of large amounts of input noise. These filters may be viewed as an extension of binary phase-only filters and similarly are motivated by the prospect of near-term real-time implementation.

I. Introduction

Binary phase-only filters (BPOFs) are well known, 1.2 and their implementation using magnetooptic spatial light modulators (SLMs)³ has been accomplished with good agreement demonstrated between experiment and theory. +: BPOFs encode modulation values of -1 and +1 usually based on thresholding in the complex plane of a reference pattern transform. Advantages of BPOF correlation include narrow correlation peaks (corresponding to intrinsic edge enhancement of the filter impulse response), high light efficiency (typically 50 times more light in a correlation peak than for the corresponding ideal matched filter), availability of a relatively well-developed SLM for implementation, and a greatly reduced information content compared with the matched filter (which allows many more filters to be stored on-line in a system).

Practical use of real-time correlation demands filter designs which maximize the ratio of (properly located) target-class response to all nontarget class responses. This, the smart filter design issue, is application dependent, and thus a single technique may not be best in all cases. Usually other performance issues constrain the design, such as the desire to maintain narrow correlation peaks and high optical throughput. Smart

filter formulations for both matched filter and BPOF encoding have been reported8-16 and have demonstrated promising correlation performance in simulations. However, BPOF filters, irrespective of formulation. have the intrinsic property of passing all the light from the input scene to the correlation plane. This leads to desirable characteristics such as narrow correlation peaks and high efficiency, but it comprises a fundamental limitation on performance in the area of rejecting nontarget inputs. It would seem that the best a BPOF can do to minimize output (correlation) plane clutter resulting from a nontarget input pattern is to distribute uniformly the associated energy over the output plane, thus minimizing the expected value at a given position. This concept was used to define the minimum average correlation energy (MACE) filter. 5 although the formulation addressed only continuousvalued filters.

The above rationale motivates the incorporation of amplitude modulation in the correlation filter. The matched filter of course has this, but we desire a filter encoding scheme supported by available SLM technology without the need for holographic carriers and their attendant space—bandwidth demands. The following sections propose such an encoding scheme, the ternary phase-amplitude filter (TPAF) and a transform-ratio formulation for designing noise immunity and discrimination into TPAFs. Simulation results demonstrating the performance of the resulting filter designs are presented.

II. Ternary Phase-Amplitude Filter and Transform-Ratio Formulation

The TPAF encoding represents the simplest possible step beyond BPOF encoding: a zero-modulation

David L. Flannery and John S. Loomis are with University of Dayton Research Institute, Applied Physics Division, Dayton, Ohio 45469-0001; Mary E. Milkovich was with University of Dayton Res. Institute when work was performed and is now with Eastman Kodak Company, Rochester, New York 14650.

Received 23 February 1988. 0003-6935/88/194079-05802.00/0. © 1988 Optical Society of America.

state is added. Thus the filter encodes the modulations -1.0, and 1. This may be conceived as a product of a binary-phase modulation pattern and a binaryamplitude modulation pattern, which could be implemented by cascading known binary SLMs including magnetooptic3 and ferroelectric liquid crystal17 devices. The magnetooptic SLMs3 are reported to have a third modulation state, which, with proper optical configuration, could furnish zero-modulation, thus implementing the TPAF in a single device. These devices have demonstrated the potential for real-time electrical addressing at attractive rates,3.17 thus enabling adaptive correlation system concepts in which thousands of reference patterns are efficiently stored on-line as simple binary patterns in random access digital memory.

The rationale for the TPAF starts with that already given. The zero-modulation allows spatial spectral components carrying mainly nontarget energy to be blocked, thus preventing that energy from reaching the correlation plane, where it could only have detrimental effects on correlation performance. The transform-ratio formulation bases the choice of spatial frequency elements to be blocked (assigned zero-modulation) on the ratio of spectral energies (power spectral densities) of the target and nontarget patterns at each spatial frequency. In the most general case, these spectra are defined statistically, i.e., the Fourier transforms of autocorrelation functions for the target and nontarget ensembles. The application to background noise models having knr : ower spectral densities is obvious. For inputs and contargets not described by statistical models. program obtained by averaging over statistically representative target and nontarget training sets would be used.

A simple type of transform-ratio TPAF (TR-TPAF), which we have investigated and report here, starts with a cosine-BPOF¹⁸ binary-phase pattern computed for a reference pattern and sets selected filter frequency elements to zero-modulation based on the transform-ratio concept discussed. We note that ternary phase-amplitude filters have been reported, ^{10,19} but in those cases zero-modulation was either incorporated in the form of a simple low spatial frequency cutoff ¹⁰ or as a result of two-level quantization of a continuous complex-valued filter of the correlation SDF type. ¹⁹ Both formulations are quite distinct from the transform-ratio approach reported here.

III. Transform-Ratio TPAF Correlation Simulations

To investigate the TR-TPAF filter, correlations were simulated using discrete Fourier transforms on a 64 × 64 sample format. The binary target and nontarget inputs used are shown in Fig. 1. The target class consisted of a single case, the gunboat of Fig. 1(a), while nontarget classes included a single case of the gunless boat of Fig. 1(b) and random binary background noise. We note that the use of binary input and reference patterns is an arbitrary choice for these simulations, not implied by the TR-TPAF formulation.




Fig. 1. Binary target and nontargets used in correlation simulations: (a) target gunboat; (b) gunless boat; (c) target superimposed on background noise of level 0.4.

Background noise was added per the following recipe: A spatially random distribution of samples outside the boats was set to amplitude one. The noise level was defined as the probability of a background element being turned on and thus approximately represents the fractional area coverage of the noise. We believe this technique closely matches a binary image

which would result from thresholding an image including an object superimposed on uniformly distributed random background clutter. The noise model was tested by observing the discrete Fourier transforms of noise samples, which exhibited tlat (white) spectra (qualitatively judged) except for a large dc value. The latter is expected because the noise is not zero-mean. Figure 1(c) is an example of the target boat superimposed on noise of level 0.4.

The design goal of the TR-TPAF in these simulations was to maximize simultaneously correlation discrimination and signal-to-clutter (S/C) performance in the presence of large amounts of input noise. Discrimination is defined as the ratio of the energies (simulated intensities, i.e., magnitude-squared) of the properly located (gunboat) target correlation peak sample to the highest value associated with the gunless boat cross-correlation pattern, all other conditions being identical. We note that correlations were not performed with both boats simultaneously present in the input scene. Rather separate correlations were performed with each boat superimposed on the same noise background, and the resulting peak values were used to form this ratio.

Signal-to-clutter (S/C) is defined as the ratio of the properly located target correlation peak to the highest peak elsewhere in the correlation plane regardless of whether it was associated with input noise or was a secondary peak associated with the target correlation pattern. This strict definition is adopted with simple threshold processing of correlation plane patterns in mind. A secondary peak, or shoulder, of a correlation pattern hampers such processing just as much as do true clutter responses.

The use of single-sample values in the definitions of S/C and discrimination is consistent with the envisioned simple threshold correlation pattern processing and is made practical by the narrow correlation peaks characteristic of BPOF and TPAF correlation, as observed in our simulations. We note that in none of the cases reported was a S/C value limited by a secondary peak of the correlation pattern.

The qualifier "properly located" is included in our definitions in keeping with the assumed requirement to use the correlation peak location determined by simple thresholding as an estimate of target location, which would be hampered if the peak location were not correctly registered with the input target reference centering point used in computing the filter. In all cases reported here, the correlation peak was properly located.

The TR-TPAFs were made by starting with the cosine-BPOF computed for the gunboat pattern, Fig. 1(a), and setting elements to zero if the gunboat transform energy at that element did not exceed each of two threshold values, one for each of the nontarget classes. The discrimination threshold was defined as a transform-ratio factor times the spectral energy of the gunless boat at the frequency element under consideration. The noise-blocking threshold was defined as a constant value expressed as a fraction of peak Fourier

the gunboat target pattern. The use of a uniform threshold is based on the flat spectrum of the back ground noise used. An additional ad hoc modification was incorporated in all the filters reported here. The dc element of the filter was always set to zero-modulation to block the huge dc component of the background noise used. This is consistent with the transform-ratio concept. Simulations run without this modification invariably yielded inferior results.

The transform-ratio factor and the noise-blocking threshold comprise two adjustable design parameters of the TR-TPAF for this case, and in general performance must be explored over the design space thus defined. Our investigations did not span the entire design space but covered a region of interest. The dependence of discrimination on the transform-ratio factor was explored initially with zero background noise in the inputs and zero noise-blocking threshold in the filters. A transform ratio of 1.21 was found to give good, although not necessarily the best, discrimination. This value was adopted for all the simulations reported here. We note that the simple cosine-BPOF filter for the gunboat provided 2.6-dB discrimination with zero input noise, whereas the TR-TPAF just discussed yielded 5.7 dB, again all with zero noise in the input scenes.

A series of simulations were performed with varying amounts of input noise and parametric variations of the TR-TPAF noise-blocking threshold. The results are plotted in Fig. 2. Only one background noise sample was used for each data point plotted; thus the values do not represent proper averages over the noise ensemble. However, the trends are readily apparent. The TR-TPAFs with appropriately set noise-blocking thresholds sustain better discrimination and S.C performance with higher levels of input noise. As a point of reference, the simple cosine-BPOF yields discrimination of 1.14 dB and S/C of 1.95 dB with an input noise level of 0.4. The absolute values are not too important, since they are subject to many arbitrary factors, such as resolution and choice of target and nontarget shapes.

The cases plotted correspond to blocking from 40 3 to 89.1% of the filter elements. In the latter case corresponding to a filter noise-blocking threshold of 0.04, the energy in the peak correlation sample using the TR-TPAF was 8.4 dB below that using the cosine-BPOF. Thus a major portion of the efficiency advantage of BPOFs relative to matched filters (typically 17 dB) has been sacrificed to improve S/C and discrimination performance.

Broadening of correlation peaks was observed for TR-TPAF correlations relative to the essentially single-sample peaks exhibited by BPOF correlations. No attempt was made to quantify the broadening, but subjective assessment of several of the worst cases indicates at most a 2:1 broadening. Although significant, this is expected to have minimal impact on most applications.

None of the TR-TPAFs used in Fig. 2 could provide




Fig. 2. Plots of signal-to-clutter (a) and discrimination (b) performance of TR-TPAF correlations.

successful target location when the input noise level was increased to 0.6, or could the simple cosine-BPOF. TR-TPAFs made with a transform-ratio factor of 0 (i.e., giving up discrimination against the gunless boat) were able to provide positive S/C ratios, thus affording target location.

An important issue for any coherent optical processing system is the effect of nonideal optical components on performance. Phase distortions in the Fourier plane (e.g., associated with an SLM) are of particular interest. To investigate this, we introduced spatially random phase errors uniformly distributed over ±45° at the Fourier plane for a simulation involving a filter noise-blocking threshold of 0.04 and an input noise level of 0.4. The simulation was repeated for seven different random phase sample sequences from the computer random number generator. The S/C ratio ranged from 2.05 dB less to 0.28 dB greater than that of the corresponding case with zero phase error and averaged 1.81 dB less. Correlation efficiency (i.e., peak energy) was reduced from 0.13 to 2.32 dB with an average reduction of 1.32 dB. The degradation of correlation performance is to be expected. The large variation over noise samples does not seem surprising.

We postulate that the phase noise alters correlation to different degrees depending on alignment of a particular noise sample with the filter function, actually increasing correlation efficiency in one case as stated. This is possible since the encoded binary phase filter values are nonoptimum compromise values.

IV. Conclusion

The TPAF filter formulated using the transformratio concept has shown improved discrimination and noise immunity in simulations of coherent optical correlation. Near-term real-time implementation should be relatively easy due to the simple modulation levels required (-1,0,+1) and is desirable to verify the simulation results. Some reduction of correlation light efficiency and broadening of correlation peaks, relative to BPOF correlation, was observed and is believed an inescapable price in trade for improvements in other performance parameters. The TR-TPAF still provides better light efficiency and narrower peaks than does classical matched filtering.

The effect of substantial (quarter wavelength peakto-peak) phase distortion in the filter plane was investigated for one case and resulted in significant degradation of average correlation performance as expected. Although not addressed in these simulations, the TR-TPAF should apply to the design of smart filters having the goal of distortion-invariant performance. This most important aspect of TPAF filter design should be addressed in the near future.

This work was supported in part by Air Force Rome Air Development Center contract F19628-87-C-0073.

References

- J. Horner and H. Bartelt, "Two-Bit Correlation," Appl. Opt. 24, 2889 (1985).
- 2 J. Horner and J. Leger, "Pattern Recognition with Binary Phase-Only Filters," Appl Opt. 24, 609 (1985).
- W. Ross, D. Psaltis, and R. Anderson, "Two-Dimensional Magneto-Optic Spatial Light Modulator for Signal Processing," Proc. Soc. Photo-Opt. Instrum. Eng. 341, 191 (1982).
- D. Psaltis, E. Paek, and S. Venkatesh, "Optical Image Correlation with a Binary Spatial Light Modulator," Opt. Eng. 23, 698 (1984).
- D. Flannery, A. Biernacki, J. Loomis, and S. Cartwright, "Real-Time Coherent Correlator Using Binary Magnetooptic Spatial Light Modulators and Input and Fourier Planes," Appl. Opt. 25, 466 (1986).
- P. Keller, D. Flannery, S. Cartwright, and J. Loomis, "Performance of Binary Phase-Only Correlation on Machine Vision Imagery," Proc. Soc. Photo-Opt. Instrum. 728, 257 (1986).
- D. Flannery, J. Loomis, M. Milkovich, and P. Keller, "Application of Binary Phase-Only Correlation to Machine Vision," Opt. Eng. 27, 309 (1988).

- 3. V. Sharma and D. Casasent. "Optical Linear Communications," Proc. Soc. Photo-Opt. Instrum. Eng. 519, 10, 1284
- R. Kallman, "Construction of Low Noise Optical Correlation Filters," Appl. Opt. 25, 1032 (1986).
- R. Kallman, "Optical Low Noise Phase-Only and Binary Phase Only Optical Correlation Filters for Threshold Detectors Appl. Opt. 25, 4216 (1986).
- R. Kallman, "Direct Construction of Phase-Only Correlation Filters," Appl. Opt. 26, 5200 (1987).
- G. Schils and D. Sweeney, "Iterative Technique for the Synthesis of Distortion-Invariant Optical Correlation Filters," OptiLett. 12, 307 (1987).
- A. Mahalanobis, B. Kumar, and D. Casasent, "Spatial-Tempral Correlation Filter for In-Plane Distortion Invariance," Apr. Opt. 25, 4466 (1986).
- D. Jared, D. Ennis, and S. Dreskin, "Evaluation of Binary Phase-Only Filters for Distortion-Invariant Pattern Rec gn. tion," Proc. Soc. Photo-Opt. Instrum. Eng. 884, 139 (1988)
- J. Rosen and J. Shamir, "Composite Phase Filters for Distortion Invariant Pattern Recognition," Proc. Soc. Photo-Opt. Instrum-Eng. 813, 285 (1987).
- 16 A. Mahalanobis, B. Kumar, and D. Casasent, "Minimum Average Correlation Energy Filters," Appl. Opt. 26, 3633 (1987)
- N. Clark and M. Handschy, "Device Application of Ferroelectric Liquid Crystals," Proc. Soc. Photo-Opt. Instrum. Eng. 625, pp. (1986)
- D. Cottrell, R. Lilly, J. Davis, and T. Day, "Optical Correlation Performance of Binary Phase-Only Filters Using Fourier and Hartley Transforms," Appl. Opt. 26, 3755 (1987).
- D. Casasent and W. Rozzi, "Computer-Generated and Phase-Only Synthetic Discriminant Function Filters," Appl. Opt. 25, 3767 (1986).

Postdoctoral Research Fellowships in Chemistry

The Division of Chemistry, in NSF's Directorate for Mathematical and Physical Sciences, plans a new round of two-year Postdoctoral Research Fellowships in Chemistry: Approximately 30 awards carrying a \$26,000 stipend to the fellow. plus institutional and supply allowances. will be awarded in March 1989. Eligibility is restricted to citizens or nationals of the U.S. receiving Ph.D. degrees between June 1, 1988, and Sept. 30, 1989. The Fellowships are tenurable at any eligible laboratory providing an outstanding research environment in an emerging research area different from the subfield in which the doctoral research was accomplished. A special feature of these Fellowships is a starter-grant option, in which any Fellow accepting a tenure-track academic appointment after the first or second year of the Fellowship may apply for a \$32,000 starter grant, contingent on receiving matching funds from the employing college or university

Applications may be obtained from the Chemistry Division's Special Projects Office (357-7503) after Oct 1, 1988. The deadline date for submission of applications will be Dec. 15, 1988.

Design elements of binary phase-only correlation filters

David L. Flannery, John S. Loomis, and Mary E. Milkovich

The basic design elements of binary phase-only filters (BPOFs) include the angle of a thresholding line in the complex plane and the offset from a center reference point of the pattern transformed in constructing the filter. These factors are analyzed, and a general formalism for the threshold-line angle variation is presented and related to the special cases of cosine-, sine-, and Hartley-BPOFs and to the general characteristics of the BPOF impulse response. The effects of these elements in correlation performance are investigated using computer simulations with realistic target and clutter patterns. The results indicate significant, but not major, variations of correlation performance for the cases studied.

I. Introduction

The binary phase-only filter (BPOF) formulation 1.2 has shown promise as a Fourier plane filter for coherent optical correlation in both computer simulations³ and experimental demonstrations. 4.5 The BPOF is a particular class of phase-only filter (POF),6.7 in which only the two-phase modulation values, zero and π rad (or equivalently the modulation values -1 and 1), are encoded. Originally the choice of encoded value at each resolvable point in the filter plane was based on thresholding on the real1 or imaginary2 value of the computed (discrete) Fourier transform of the desired reference pattern. These are designated cosine- or sine-BPOFs, respectively,8 as they can be derived from cosine or sine transforms. Due to fundamental symmetry properties of the Fourier transform, these BPOFs encode information derived from only the even or odd components of a real reference pattern, the usuai case.

Horner and Leger² showed that the sine-BPOF impulse response approximately comprised a superposition of two edge-enhanced replicas of the original reference shape, one aligned as in the reference scene and one rotated 180°. The same behavior is observed for cosine-BPOFs. The edge enhancement is an expected characteristic of any phase-only filter. Correlation with the 180° rotated component amounts to convolution and fortunately results in little degradation of the desired correlation peak in many cases of interest.

Cottrell et al.⁸ investigated the effects of the rotated component for the sine-, cosine-, and Hartley-BPOFs and found that the Hartley-BPOF was superior for correlation in the cases they studied.

For even or odd functions, the phase of the transform has only two possible values which differ by π . Thus it may easily be verified that the cosine- or sine-BPOF encodes the exact phase (within an additive physically meaningless constant) for those two cases, respectively, and thus comprises a true POF for those special cases. An even function can by synthesized by adding two replicas of an arbitrary function, one of which is rotated 180°. These coincide with the unrotated and rotated components of the BPOF impulse response.

In this paper we provide a unifying formalism in which the coame-, sine-, and Hartley transforms are particular points on a continuum of the angle of a complex-plane thresholding line used to generate the BPOF from the reference transform, and we analyze the relationship of Fourier, input plane, and impulse response variations associated with variations of this angle. We note that variations of the thresholding-line angle have already been used to optimize BPOFs by Kallman.⁹ To provide a practical assessment of the effects of this design parameter we have performed correlation simulations using realistic imagery on a 256 × 256 sample format. These results are presented and interpreted.

II. BPOF Basic Design Elements

We consider the effects of two design parameters: reference pattern centering and thresholding-line angle (TLA).

A. Centering Effects

Centering refers to offset of the reference pattern relative to the central axis of the input plane. The

When this work was done all authors were with University of Dayton Research Institute, Applied Physics Division, Dayton, Ohio 45469; M. E. Milkovich is now with Eastman Kodak Co., Rochester, New York 14600.

Received 23 February 1988. 0003-6935/88/204231-05802.00/0. © 1988 Optical Society of America.




Fig. 1. Conventions defining thresholding line angle in complex plane.

impulse response of the resulting (computed) BPOF exhibits two 180° rotated replicas, and the relative positioning of these is given by the centering offset chosen.2 We refer to the rotated components as the correlation (presumably desired) and convolution (presumably undesired) components. The relative positions of these components determine how the correlation and convolution responses overlap. This overlap can enhance or degrade important performance parameters such as correlation peak value and the ratio of correlation peak to other correlation plane values (called clutter responses), as has been pointed out.8 The effects are strongly dependent on the nature of the particular reference and input patterns involved. Thus the best design procedure is probably to try to anticipate these effects based on an intuitive understanding, as just outlined, and to choose the input pattern offset accordingly. Several design iterations using a computer model to simulate correlation performance with varying offsets would be appropriate. Correlation peak location is shifted in accordance with the filter shift design parameter, and this offset must be treated in interpreting the correlation result for location estimation.2 The required correction is a known translation, but the situation becomes complicated if the designer wishes to use the convolution response to recognize 180° rotations of the target pattern, since the shift correction is opposite for that response, and there would appear to be no way to ascertain which version of the target pattern led to an observed correlation-plane peak.

B. Threshold Angle Effects

Figure 1 indicates our choice of convention for defining the TLA in the complex plane. Thresholding a reference pattern's Fourier transform based on this line generates the BPOF as follows:

$$F_{r}\cos(\theta) - F_{r}\sin(\theta) \begin{cases} \geq 0; & \text{BPOF} = 1, \\ < 0; & \text{BPOF} = -1, \end{cases}$$
 (1)

where F_r and F_t are the real and imaginary parts of a Fourier transform value.

The choice of 0, 45, and 90° for θ may easily be verified to correspond to the cosine-, Hartley-, and sine-BPOFs, respectively.⁸ (In the 90° case, thresholding is on the negative of the imaginary part of the

transform instead of the positive value normally used in the sine-BPOF definition, a trivial difference.

Now consider a real 2-D pattern $f(\bar{x})$ with Fourier transform $F(\bar{k})$. We know that the transform $f(-\bar{x})$ rotated 180°) is $F^*(\bar{k})$. If we synthesize a new pattern consisting of a superposition of 180° rotated and phase-shifted versions of $f(\bar{x})$, i.e.,

$$f'(\bar{x}) = \exp^{-z} f(\bar{x}) + \exp^{-z} f(-\bar{x})$$

the corresponding Fourier transform is

$$F'(\bar{k}) = \exp^{-2}F(\bar{k}) + \exp^{-2}F^*(\bar{k}).$$

This can be simplified to

$$F' = 2[F, \cos(\phi) - F, \sin(\phi)],$$

where the dependence of \bar{k} has been made implicit. We note that this transform is real, and the cosine-BPOF for this transform would be determined by thresholding on F' [Eq. (4)].

Now compare Eq. (1) with Eq. (4). If we make the correspondence $\theta = \phi$, exactly the same BPOF pattern is being defined. Thus we conclude that the BPOF for a given pattern made using a TLA of θ is identical to the POF for a pattern comprised of the sum of the original pattern and its 180°-rotated replica superimposed with $\pm \theta$ relative phase shift.

This analysis is valuable in two ways. First, we see that the previously defined BPOF types (sine, cosine. Hartley) may be viewed as three points on the continuum of TLAs applied in constructing the BPOF. There is no fundamental or practical reason to restrict BPOF designs to one or more of these particular points. Second, the relationship shown between the BPOF made with a particular TLA and a POF made for a synthesized input of 180°-rotated and phase-shifted components suggests an intuitive approach to predicting the impulse response of the BPOF.

We know that the impulse response of a POF is an edge-enhanced replica of the reference pattern used to construct the POF. Applying this concept allows us to predict intuitively the impulse response of a BPOF (made with a particular TLA). We would expect to see 180°-rotated edge-enhanced replicas of the reference pattern, and in those areas where the replicas exhibit little overlap we would expect them to have phases equal to ±TLA. The latter prediction is purely intuitive. The nonlinear operations implicit in POF generation preclude a simple analysis to verify it. However, we have investigated this effect experimentally.

Figure 2 shows BPOF impulse responses for a tank image used later in our studies. The magnitude-squared of the impulse responses was thresholded at an arbitrarily chosen level and plotted. BPOFs and impulse responses for TLAs of 22.5, 45, 67.5, and 90° were computed. The figure shows two of these cases corresponding to two different centering offsets and TLAs. The effects of TLA variations, primarily exhibited in phase, are not pronounced in magnitude-squared plots; however, the figure is marked to indicate some of the locations where the phase of the two replicas was computed. The measured phase varied




Fig. 2. BPOF impulse response of tank target pattern for two different centering offsets. The arrows indicate some of the points where phase was computed to explore its variation with TLA.

from the intuitive prediction (±TLA) by only 1.86° rms (4.7° in the worst case) over the four TLA cases. The standard deviation of measurements for a particular TLA case varied from 0.08 to 2.2°. We believe these thirty-six phase measurements involving four TLA values offer strong support for the intuitive premise.

The enhanced ability to predict impulse responses, including phase in nonoverlapping areas, afforded by the above analysis and empirical observations, should aid the design of BPOFs with regard to centering and TLA parameters.

M. Parametric Study of Design Element Effects

To study the effects of the two basic BPOF design elements, we computed simulated BPOF correlations using FFT algorithms on a 256 × 256 sample format using actual tank images superimposed on actual terrain backgrounds. A total of forty-six correlations were simulated to investigate parametric variations of centering offset and TLA on correlation performance factors including signal-to-clutter (S/C), discrimination, and Horner efficiency.¹⁰

The performance terms are defined:



Fig. 3. Typical correlation plane result, plotted linear in intensity



Fig. 4. Input scene for first series of simulations showing target tank and background clutter.

Discrimination is the ratio (usually expressed in decibels) of the energy (magnitude-squared) of the desired (target) tank's correlation peak sample to the highest such sample value located within the outline of an undesired (nontarget) tank. (This metric applies only in cases including a nontarget tank in the input.)

The S/C ratio is the ratio of the desired correlation peak energy to the next highest sample value anywhere in the correlation plane, excluding those assigned to a nontarget tank, if present.

Horner efficiency is defined in a simple modified sense as the fraction of the energy incident on the filter plane which is present in the single desired correlation peak sample. This simplified definition is reasonable because these correlations, as do POF and BPOF correlations, in general, exhibit peaks which are essentially only one sample wide. Figure 3 is a typical correlation output using a vertical scale, which is linear in intensity (magnitude-squared).

In a first series of simulations, only a single target tank was used (see Fig. 4), and the emphasis was on S₁C performance. A second series used a different terrain background, and a second tank of a different type was added to enable discrimination to be measured (see Fig. 5). In all eases the isolated image of the tank in Fig. 4 was the target, and BPOFs were generated using its transform.

A. Signal-to-Clutter Variations

Figure 6 is a scatter plot of the S/C ratios for the first series of simulations. The plot summarizes the effects of both centering (five different shifts) and TLA (five angles spanning 0-90°). Note that TLA effects are



Fig. 5. Input scene for second series of correlation simulations.

theoretically symmetric about TLA = 0°. Changing the sign of the TLA causes only a uniform phase change in the resulting correlation function, and this is of no practical consequence. Thus 0-90° is a maximum nonredundant range for TLA. This was verified by several correlations performed with negative TLAs. which yielded results exactly identical (in intensity) to their positive TLA counterparts. The centering shifts are expressed in numbers of samples on the 256 \times 256 sample format using an east-west-north-south (E-W-N-S) convention.

First, we note that the entire data set exhibits only a minor dependence on either centering or TLA; the standard deviation of S/C over all the data is only 0.5 dB compared to an overall average value of 7.72 dB. Averaged over all centering shifts, the values range from 7.33 to 8.30 dB as a function of TLA. Averaged over all TLAs the variation with centering shift is even less: 7.59-7.89 dB. In spite of the above, the eye does tend to perceive a systematic trend in which S/C is highest for a TLA of 67.5°. We suspect that this is particular to the specific patterns chosen for correlation rather than a general result. If the result is taken to have any significance, it is interesting to note that best performance occurs for a TLA differing from any of the previously defined BPOF types. The TLA of 67.5° is midway between the Hartley- and sine-BPOF

Considering the relatively small standard deviation of all the S/C values the only conclusion supported by these data is that significant, but not major, variations occur with TLA, and no significant variation with centering is exhibited. The TLA variation may be optimized near 67.5° for the specific pattern correlated.

Discrimination Variations

Figure 7 plots the discrimination results computed from the second series of simulations using the input scene shown in Fig. 5. These results seem more scattered than the S/C results, yet the standard deviation over the entire set is still relatively low: 0.65 dB compared to an overall average discrimination of 5.56 dB. Variation with TLA (averaged over centering shift) is from 4.83 to 6.03 dB, the latter corresponding to 67.5°. The spread of values is noticeably less at 45 and 67.5°





Fig. 6. Scatter plot of signal-to-clutter ratio (dB) for first series Centering shifts are designated by N-S-E-W directions and are measured in units of samples on a 256 × 256 format



-N15, W25 • 520 Fig. 7. Discrimination (dB) between tanks measured for second series. Centering shifts are designated by N- E-W directions and are measured in units of samples on a 256 × 256 format.

4E32

than at 0 and 90°. Variation with centering shift (averaged over TLA) is again small: 5.45-5.80 dB. We know no explanation for either the best average discrimination for a TLA of 67.5° (which was only slightly better than at 45°) or for the reduced data spread for Hartley and 67.5° cases compared with the sine-BPOF and cosine-BPOF cases. We note that the latter effect may be related to the mixture of odd and even target pattern information encoded for TLA angles other than the cosine and sine BPOF cases.

Input Shift Effects

Although not a BPOF design issue, we were curious as to the effect on discrimination of shifting the target tank location in the input scene (see Fig. 5) The target tank was shifted fifteen samples north and fifteen samples west, and correlations were simulated with BPOFs made with zero centering shift and for four values of TLA: 0, 22.5, 45, and 67.5°. The results were quite similar to those of the corresponding simulations with the tank not shifted. Average discrimination was 5.14 dB compared with 5.56 dB for the unshifted cases.

n Homer Efficiency

The Horner efficiency ranged from 0.011 to 0.026% over all the simulations performed, with no systematic dependence on filter design parameters discernible. Since the BPOF passes all the light incident in the input pattern Fourier transform, the Horner efficiency also gives the fraction of input scene energy, which is concentrated in the correlation peak sample. For these scenes the average energy per sample was about the same for the tank target and backgrounds. Based on the size of the tank, ~5% of the total input scene energy was included in the tank. Thus the observed Horner efficiencies correspond to 0.22-0.52% of input target energy appearing in the correlation peak samiple.

IV. Conclusions

The effects of two basic BPOF design elements on correlation performance using realistic high-resolution target-in-background imaginary have been assessed, with the result that no variations of major significance were found for the particular correlation cases used. Apparent slightly superior correlation performance was exhibited for TLA angles intermediate to the sine- and cosine-BPOF cases. We conclude that these results exhibit insufficient systematic dependencies of performance on TLA or centering to furnish a basis for generalization as to the choice of an optimum TLA angle. We speculate that scene-specific relative spatial frequency phase alignments of the target and nontarget patterns combine with the oddeven sensitivity variation of BPOFs with TLA to yield these results. Thus the optimum choice of TLA must be investigated for each correlation application.

Analysis shows that the previously defined BPOF filter types (sine, cosine, and Hartley) are encompassed as special cases in the variation of one of the basic BPOF design parameters, the thresholding-line angle (TLA), applied in constructing a BPOF from the Fourier transform of a reference pattern. There is no intrinsic justification for restricting filter designs to the three special cases, and performance simulations showed a slight preference for a TLA between the

Hartley and sine points in some cases. After initial submission of this paper we learned of an alternate and apparently equivalent formalism for treating TLA variations called the generalized Hartley filter.

Centering shift was the other design parameter studied, and our results show insignificant dependence of correlation performance on it for the cases studied.

Analysis yielded a simple interpretation of the impulse response of BPOFs made with various TLAs in terms of POFs made with inputs synthesized by superimposing two 180°-rotated and phase-shifted replicas of the original input function. This provides intuitive understanding helpful in predicting the impulse response of the filters, and experimental results verified these intuitive expectations.

References

- D. Psaltis, E. Paek, and S. Venkatesh, "Optical Image Correlation with a Binary Spatial Light Modulator," Opt. Eng. 23, 625 (1984).
- J. Horner and J. Leger, "Pattern Recognition with Binary Phase-Only Filters," Appl. Opt. 24, 609 (1985).
- J. Horner and H. Bartelt, "Two-Bit Correlation." Appl. Opt. 24, 2889 (1985).
- D. Flannery, A. Biernacki, J. Loomis, and S. Cartwright. 'Real Time Coherent Correlator Using Binary Magnetooptic Spatial Light Modulators and Input and Fourier Planes.' Appl. Opt. 25, 466 (1986).
- P. Keller, D. Flannery, S. Cartwright, and J. Loomis, "Performance of Binary Phase-Only Correlation on Machine Vision Imagery," Proc. Soc. Photo-Opt. Instrum. Eng. 728, 257 (Oct. 1986).
- J. Horner and P. Gianino, "Phase-Only Matched Filtering. Appl. Opt. 23, 812 (1984).
- P. Gianino and J. Horner, "Additional Properties of the Phase-Only Correlation Filter," Opt. Eng. 23, 695 (1984).
- D. Cottrell, R. Lilly, J. Davis, and T. Day, "Optical Correlator Performance of Binary Phase-Only Filters Using Fourier and Hartley Transforms," Appl. Opt. 26, 3755 (1987).
- R. Kallman, "Optimal Low Noise Phase-Only and Binary Phase-Only Optical Correlation Filters for Threshold Detectors," Appl. Opt. 25, 4216 (1986).
- J. Horner, "Light Utilization in Optical Correlators." Appl. Opt. 21, 4511 (1982).
- F. Dickey, J. Mason, and K. Stalker, "Analysis of Binarized Hartley Phase-Only Filter Performance with Respect to Stochastic Noise," Proc. Soc. Photo-Opt. Instrum. Eng. 983, 266 (1988).

This work was supported by Air Force Rome Air Development Center contract F19628-87-C-0073.

APPENDIX D

New Formulations for Discrete-Valued Correlation Filters

David Flannery, John Loomis, and Mary Milkovich The University of Dayton Research Institute Dayton, Ohio 45469

ABSTRACT

Ternary correlation filters, which encode modulations of -1, 0 and +1, may be viewed as a logical step beyond binary phase-only filters. Both formulations are motivated by the prospect of relatively simple real-time implementation compared to full complex matched filters. The sero-modulation state of ternary filters affords additional flexibility and control in filter design. In particular both enhanced nontarget discrimination and reduced distortion sensitivity, relative to simple binary phase-only filters, can be achieved simultaneously as demonstrated by correlation simulations reported here.

1. INTRODUCTION

Binary phase-only filters (BPOF) are well known in theory and practice [1,2]. They offer good basic correlation performance including narrow peaks, high efficiency, and good discrimination due to intrinsic high frequency emphasis relative to matched filters. Practical advantages of the BPOF include reduced information content, allowing efficient digital storage of many filters, and availability of a spatial light modulator (SLM) well suited to real-time implementation of the filter [2,3].

Practical application of real-time correlation demands filter designs which provide distortion invariance while maintaining sufficient discrimination against nontarget patterns; this defines the "smart filter" problem. Smart filters, such as synthetic discrimination function (SDF) filters [4], have been defined for continuous-valued encoding using linear algebraic formulations, but these approaches are not rigorously applicable for BPOF formulation due to the implicit nonlinearity involved in BPOF generation. Recently the filter-SDF (ISDF) formulation [5,6], based on a simple iterative solution, was reported and appears to provide a viable BPOF version of the SDF. Correlation peak responses can be forced to take on a specified pattern over a set of training images (although convergence of the iterative procedure is not guaranteed for any arbitrary pattern), thus affording a technique for designing distortion invariance and out-of-class rejection into a BPOF. Kallman [7,8] also has reported impressive distortion-invariant BPOFs with good clutter rejection, although his recipe appears to be-computationally intensive compared with the ISDF-BPOF and techniques to be reported here.

We have previously reported the ternary phase-amplitude filter (TPAF) concept and initial simulations using a "transform-ratio" technique which demonstrated enhanced discrimination against noise and nontarget patterns of To review, a ternary filter encodes the modulation values -1, 0, and +1, i.e. it implements a sero-modulation state in addition to those of the BPOF. The rationale for expecting improved correlation performance as a result of intelligent use of the sero-modulation state is as follows: A BPOF passes all the Fourier transform energy from the input scene since its modulation has constant unity magnitude. Thus both desired (target) as well as undesured (nontarget) energy is passed to the correlation plane where the nontarget energy can have only detrimental effects on correlation performance. This sets fundamental limits on the performance of any BPOF filter, smart or otherwise Typically, target and nontarget transforms exhibit spectral energy concentration patterns which are distinct. The sero-modulation state can be used to block energy in those regions of the spatial spectrum which contain higher ratios of nontarget to target energy, which should improve correlation performance.

The transform-ratio TPAF (TR-TPAF) previously reported [9] is a simple embodiment of the above concept. The starting point is a BPOF computed for a target pattern. Then zero-modulation is set for selected frequency plane elements based on the ratio of spatial Fourier spectral energies of the target to that of one or more non-target patterns. In the previous work two nontargets were defined: white noise and a specific nontarget object. Two transform-ratio thresholds (one for each nontarget type) are thus involved in filter formulation and serve as adjustable parameters to optimise the filter for an application.

Simulations reported [9] indicated significantly better noise and nontarget discrimination for the TR-TPAF compared with a simple BPOF for the patterns studied. However, that TPAF formulation addressed only improved

discrimination. Distortion invariance was not addressed, since the initial BPOF pattern was made for a simple single view of the target object. We have developed TPAF formulations which combine the ideas used in the TR-TPAF and the fSDF-BPOF formulations to yield both improved discrimination and reduced distortion sensitivity. These formulations will be reported here along with simulation results demonstrating their efficacy.

2. FILTER-TERNARY PHASE-AMPLITUDE FILTER FORMULATION

The fSDF-BPOF [5,6] is a BPOF generated by thresholding on a composite Fourier transform consisting of a weighted sum of training set transforms. In the reported iterative formulation, the weighting coefficients are simply adjusted until the correlation peak intensity pattern across the training set matches the prescribed pattern within a specified tolerance. In each iteration the Ad Hoc iterative procedure involves adjusting each weight in proportion to the amount by which the correlation peak intensity for the corresponding training set component differs from the prescribed pattern. Iteration is stopped when the correlation peak responses across the training set match the desired pattern within the specified tolerance.

Our new ternary filter formulation, which we have labeled the "fTPAF", combines the fSDF-BPOF idea just described with the transform-ratio ternary filter concept (TR-TPAF). First a zero-modulation pattern (essentially a binary magnitude mask) is formed using the transform-ratio technique already described. Then the identical fSDF-BPOF iteration technique is used to force equal correlation peak response over a training set. The difference is that the (ternary) filter pattern is defined as the point-by-point product of the BPOF-thresholded composite transform and the (fixed) binary magnitude mask, (and this filter is used in evaluating correlation peak responses at the appropriate point in the iterative cycle).

The mathematical description of the fTPAF is:

$$fTPAF(f_x, f_y) = BPOF\{\sum_n c_n F_n(f_x, f_y)\} \ M(f_x, f_y),$$

where f_x and f_y are spatial frequencies, n denotes a training set component, c_n is a weighting coefficient, F_n is a spatial Fourier transform, and M is the transform-ratio binary magnitude pattern. BPOF{} denotes the operation of generating BPOF values by thresholding in the complex plane of the transform value.

Note that the mathematical cascading of the binary magnitude and phase components has a direct physical counterpart – the TPAF could be implemented by cascading binary-amplitude and binary-phase modulating SLM's, if practically preferable.

3. SIMULATIONS OF STPAF PERFORMANCE

Figure 1 shows binary target and nontarget shapes used in a first series of computer simulations performed using discrete Fourier transforms over 64-by-64 sample arrays. The upper image of a crescent wrench was taken as the target. The lower pattern is the binary image of an adjustable wrench. The apparent widening of its handle near the end is exaggerated by the conversion process which produces the binary image from the original continuous tone image. We investigated correlation performance of several filter formulations in terms of both discrimination (target-vs.-nontarget) and sensitivity of target correlation peak intensity to in-plane rotations of the target. Rotations were performed digitally in an energy-preserving manner which resulted in non-binary images. The initial use of binary images was incidental, having no relationship to the types of filters under investigation.

Four types of filters were investigated: simple BPOF, TR-TPAF, (SDF-BPOF, and (TPAF). The simple BPOF filter was formed by thresholding on the real part of the Fourier transform of the target object pattern, comprising a cosine-BPOF.

A binary magnitude transform-ratio mask pattern was formed by comparing the spectral energies of the target and nontarget wrench images using a threshold ratio of 0.4, i.e., a spatial frequency element was set to zero-modulation if nontarget energy exceeded 0.4 times target energy. This resulted in 65% of the elements being set to zero. This mask pattern was point-by-point multiplied times the simple BPOF pattern to yield a TR-TPAF filter.

An (SDF-BPOF filter was formulated using a training set of five in-plane rotations (-4, -2, 0, +2, and -4 degrees) of the target wrench and forcing convergence to ±10% uniformity of correlation peak intensity over the set. The choice of two-degree increments was based on the observed rotation sensitivity of the simple BPOF filter which was ±2.2 degrees for a 3 dB drop in correlation peak intensity. Convergence required 11 iterations.




Figure 1: Inputs for first series of correlation simulations.

An ITPAF was formed using the same binary magnitude mask defined in constructing the TR-TPAF and the same in-plane rotation training set used for the ISDF-BPOF. Convergence to ±10% uniformity was obtained in six iterations.

Correlation plots, linear in intensity, for these four filters with the input image shown in Figure 1 are presented in Figure 2. The filters were compared quantitatively on the basis of three metrics: rotation sensitivity, discrimination, and correlation efficiency. These metrics must be defined.

Rotation sensitivity is defined in terms of the angular span over which target correlation peak is within 3 dB of its maximum value (expected to be at sero-degree rotation). Discrimination is the ratio of the target correlation peak intensity (magnitude-squared) to the highest sample intensity associated with the nontarget cross-correlation pattern. Correlation efficiency is defined as the ratio of the target correlation peak sample energy to the total target image energy in the input, i.e., the fraction of target image energy which is delivered to the correlation peak sample. This definition is meaningful for the correlations reported here which exhibit the narrow, near diffraction-limited, peaks characteristic of BPOF correlations.

Table 1 summarises the performance of the four filters. The TR-TPAF filter provides the highest discrimination but has about the same rotation sensitivity as the BPOF. The SDF-BPOF has decreased rotation sensitivity, by at least a factor of three relative to the BPOF, but with a sacrifice of several dB of discrimination. The fTPAF has about the same reduction in rotation sensitivity as the SDF-BPOF and also provides better discrimination than the BPOF. In cases where a range of values is given, different values were obtained depending on rotations of both the target and nontarget over a ±4 degree range, which was done to more thoroughly test the filters. The rotation range of the SDF-BPOF and fTPAF may have been more than 12 degrees, as testing was done only out to ±6 degrees.

The secondary target and nontarget peaks located to the right of the main responses are believed to be an artifact of BPOFs and TPAFs resulting from the intrinsic superimposed 180-degree rotated impulse response components combined with choice of centering of the target when making the filter. The target centering was near the left (larger) end in this case so the sub-peak results at a shift for which the 180-degree rotated component achieves maximum overlap with the input target.

We did not attempt an ISDF-BPOF trained both to reduce rotation sensitivity and reject the nontarget wrench. This should be done for the sake of completeness, since that filter would address both distortion-invariance and discrimination goals as does the ITPAF. We speculate the ITPAF has fundamental advantages due to the zero-modulation state but cannot claim this until a direct comparison is available.

The correlation efficiency is degraded by any of the smart filter formulations, and is generally poorer for the ternary filters. This is a natural result of blocking a large portion of the filter plane through sero-modulation




Figure 2: Correlation plots, (a) BPOF, (b) TR-TPAF, (c) fSDF-BPOF, (d) fTPAF.

Eller		Effection (%)	Rotation Sanathity (daggers)
	7.9	u	4.4
ecr-eror	4.1	5.5	>12 (-2.0 .030 0)
TRIPM	142	12	3.0
TPAF	65102	>04	> 12
			(·2.2 .000)

Table 1: Summary of filter performance for first series of simulations.




Figure 3: Inputs for second series of correlation simulations.

Eiler	Cinc (dill)	Eligenosis	Actesion Senestroy (degrees)	
BPGF	1.5	4.5	4.4	
(TPAF	4276	0.024-0.049	12	

Table 2: Summary of filter performance for second series of simulations.

and must be viewed as the price paid for improved overall discrimination and distortion invariance. The practical implications of the lowered efficiency for experimental implementation remain to be assessed.

The case just discussed was not a challenging one in terms of discrimination since even the simple BPOF exhibited ample discrimination. Therefore a more difficult case was contrived using the same target wrench with a modified nontarget. The new nontarget wrench was identical to the target wrench except for one end; see Figure 3. Only the BPOF and fTPAF filters were formulated for this case. Table 2 summarises their performance. The fTPAF was able to increase discrimination significantly over the marginal value exhibited by the BPOF while also reducing rotation sensitivity by a large factor. However a transform-ratio of 0.5 was necessary with attendant blocking of 88% of the Fourier plane area. The correlation efficiency of this fTPAF is extremely low, raising serious questions concerning its prospects for practical implementation. It would appear that we have asked too much of a single filter in terms of combined discrimination and distortion-invariance for this difficult case.

4. CONCLUSION

A new ternary correlation filter formulation, the ITPAF, which combines the discrimination-enhancement mechanism of transform-ratio ternary filters and the distortion-invariance approach afforded by the ISDF-BPOF, has been defined. Computer correlation simulations have supported the viability of the new formulation, including

its capability to provide combined discrimination and distortion-invariance superior to several other formulations tested. More extensive tests involving a wider variety of subject patterns and other competing filter formulations are required for a definitive assessment. Nevertheless the fTPAF appears to comprise an attractive filter design option and should be amenable to practical real-time implementation because it involves only three discrete modulation levels.

5. ACKNOWLEDGMENT

This work was supported by Air Force Rome Air Development Center contract F 19628-87-C-0073.

6. REFERENCES

- 1. J. Horner and H. Bartelt, "Two-bit correlation," Applied Optics 24, 2889 (1985).
- 2. P. Keller, D. Flannery, S. Cartwright, and J. Loomis, "Performance of Binary Phase-Only Correlation on Machine Vision Imagery," Proc. SPIE 728, 257 (1986).
- 3. W. Ross, D. Psaltis, and R. Anderson, "Two-dimensional magneto-optic spatial light modulator for signal processing," Proc. SPIE 341, 191 (1982).
- 4. V. Sharma and D. Casasent, "Optimal Linear Discriminant Functions," Proc. SPIE 519, 50 (1984).
- 5. D. Jared and D. Ennis, "Inclusion of Filter Modulation in Synthetic-Discriminant-Function Construction," Paper MA-7, Optical Society of America Annual Meeting, Rochester, NY (October 1987).
- 6. D. Jared, D. Ennis, and S. Dreskin, "Evaluation of binary phase-only filters for distortion-invariant pattern recognition," Proc. SPIE 884, (1988, to be published).
- 7. R. Kallman, "Optimal low noise phase-only and binary phase-only optical correlation filters for threshold detectors," Applied Optics 25, 4216 (1986).
- 8. R. Kallman, "Direct construction of phase-only correlation filters," Applied Optics 26, 5200 (1987).
- 9. D. Flannery and S. Cartwright, "Optical Adaptive Correlator," Proc. Third Annual Aerospace Applications of Artificial Intelligence Conference, Dayton, OH (October 1987).

Transform-ratio ternary phase-amplitude filter formulations for character recognition

Mary E. Milkovich," MEMBER SPIE David L. Flannery John S. Loomis, MEMBER SPIE University of Dayton Research Institute Dayton, Ohio 45469 Abstract. Five types of discrete-valued correlation filters, including binary phase-only and ternary phase-amplitude types, were tested in simulations addressing specific character recognition problems. The filters were evaluated for classification accuracy and correlation efficiency. Transform-ratio filter ternary phase-amplitude filters provided the best classification performance but also the lowest correlation efficiency.

Subject terms: optical signal processing; ternary phase-amplitude filter: correlation; character recognition.

Optical Engineering 28(5), 487-493 (May 1989).

CONTENTS

- 1. Introduction
- 2. Background
- 3. Filter types investigated
- 4. Simulation results
 - 4.1. Character imagery and computer software
 - 4.2. Definitions of filter performance metrics
 - 4.3. Character recognition studies
- 5. Conclusion
- 6. Acknowledgment
- 7. References

1. INTRODUCTION

This paper investigates, by simulations, the performance of several discrete-valued correlation filter formulations applied to a particular character recognition (discrimination) problem. Such filters, which encode only two or three discrete phase and/or amplitude modulation values, are motivated by considerations of near-term practical implementation. Real-time correlators using such filters show promise for both military (e.g., automatic target recognition) and commercial (e.g., robotic vision) applications.

2. BACKGROUND

The binary phase-only filter (BPOF) is a discrete-valued filter that has shown promising performance in both computer simulations and experiments. ¹⁻⁴ The BPOF exhibits desired correlation characteristics such as narrow correlation peaks and high Horner efficiencies. ³ Unlike complex valued filters, BPOFs can be implemented in a real-time experimental correlator using existing programmable spatial light modulators (SLMs), such as the Litton/Semetex magneto-optic SLM. ³ Since the BPOFs are binary (encoding the values 1 and −1), several thousands of filters can be stored in digital memory and rapidly sent out to the elec-

trically programmable SLM, making the use of these filters very promising for real-time applications.

It is well known that correlation filters, especially phaseonly filters, created from a single reference image are very sensitive to distortions of the input image relative to the reference image. The "smart filter" is a filter designed to achieve overall superior performance over a singlereference-image filter in the areas of distortion invariance and discrimination. Several smart filters have been designed for complex spatial filtering, such as lock-and-tumbler,' spatial-temporal correlation, and other linear composite filters. 411 The work of Kumar et al. 12 illustrates the fundamental trade-off between distortion invariance and discrimination for composite correlation filters. The smart filter design issue is to create a filter that will optimize this trade-off. BPOFs created from the phase of some of these complex-valued filter functions have shown improved performance over a simple BPOF filter.13 Smart filter formulations specifically created for the POF and BPOF recently have been developed.14-14

The BPOF, regardless of its formulation, passes all of the energy at the input plane of the correlator (includingnontarget energy) to the output plane. The best that the BPOF can do at reducing the correlation response to a nontarget input is to evenly distribute the corresponding energy across the correlation plane. This places a fundamental limit on the signal-to-noise ratio that this filter can achieve.

Recently, the discrete-valued transform-ratio ternary phase-amplitude filter (TR-TPAF) was developed to address this problem.¹⁷ The TPAF is a simple extension of the BPOF that encodes the filter element values -1, 0, and 1. The TR-TPAF is based on the concept that the zero-modulation state will be used to block spatial frequencies that carry mainly nontarget energy and would otherwise contribute unfavorably to the desired correlation response. Initial investigations of the TR-TPAF have shown good correlation results.¹⁷ Most important, near-term experimental implementation of the ternary filter is promising. The TPAF can be thought of as the simple application of a binary-amplitude mask to the BPOF. This can be realized experimentally by cascading two SLMs, one operating in the binary phase-only mode and the other in the binary-

^{*}Present affiliation: Eastman Kodak Co., Federal Systems Division, Rochester, N.Y.

Invited Paper IA-106 received Nov. 15, 1968; revised manuscript received Feb. 8, 1969; accepted for publication Feb. 10, 1969.
© 1969 Society of Photo-Optical Instrumentation Engineers.

amplitude mode. Both the magneto-optic¹ and the ferroelectric liquid crystal¹⁸ devices are capable of binary-amplitude modulation. The magneto-optic SLM has a third modulation state, ² which may provide the zero-modulation state.

3. FILTER TYPES INVESTIGATED

Five types of discrete-valued correlation filters were investigated in these studies: the BPOF, the TR-TPAF, the projection synthetic discriminant function BPOF (pSDF-BPOF), the filter synthetic discriminant function BPOF (fSDF-BPOF), and the filter ternary phase-amplitude filter (fTPAF). The BPOF has two design parameters that may be used to optimize performance, the centering offset and the threshold line angle (TLA) used in the Fourier plane to set the binary phase states. 19.29 For all filters investigated in these studies the reference images were centered on-axis in the reference image plane (zero offset), and the TLA was 45°, defining a Hartley-BPOF. 21

To create a TR-TPAF, a binary-amplitude mask was applied to a simple BPOF. The binary-amplitude mask was created using the transform-ratio technique.¹⁷ The zero-modulation state was set by thresholding on the ratio of the statistically defined power spectra of n in-class and m out-of-class images at each spatial frequency. The elements of the mask were set as follows: if the power spectral density of the in-class ensemble exceeded that of the out-of-class ensemble by a specified ratio, the mask element was assigned the value of unity (corresponding to a filter-element value of -1 or 1 according to the BPOF formulation); otherwise, it was set to zero.

The third filter studied was a pSDF-BPOF. A pSDF-BPOF is created from a complex-valued pSDF, or equal-correlation-peak SDF, filter. The creation of a pSDF-BPOF from a complex-valued pSDF results in a different filter function; hence, this filter is not expected to meet the goals set for the original pSDF filter.

The fourth filter studied was the fSDF-BPOF.¹⁵ This filter, unlike the pSDF-BPOF, includes the BPOF modulation in its formulation and uses iterative techniques to solve nonlinear equations for weighting coefficients of components of a composite transform, from which the BPOF is generated. Initial investigations of this filter have shown promising results.¹⁶ A major limitation of this formulation is that the iterative solution may not converge.

The fifth type of filter investigated was the TR-fTPAF. This filter formulation is similar to the fSDF-BPOF except a constant binary-amplitude mask (formulated using the transform-ratio technique) is incorporated into the nonlinear equation, which is solved iteratively for the weighting coefficients that determine the binary phase pattern.²²

It should be noted that all of these filters provide space-invariant performance. The smart filters require a training set of images representing expected in-class distortions and at least a statistical (power spectrum) knowledge of the expected out-of-class patterns. These constraints are typical of smart filter formulations for class discrimination. *11

4. SIMULATION RESULTS

A specific character recognition problem was defined and simulations of optical correlation were performed to

E	E	F
HEW YORK (NY-E)	TIFFANY (TF-E)	NY-P
E	E	R
MENERIS (NE-E)	GENEVA (GE-E)	MY-R
t	E	L
COPENHAGEN (CP-E)	MANHATTAN (MN-E)	4X-T
E	R	P
STARPLEET (SF-E)	PIERCE (PI-E)	мү-р
E	E	B
STUTTGART (ST-E)	SAIGON (SG-E)	NY-B
E	E	8
COURTER (CD-E)	SAN DIEGO (SD-E)	NY-6
E	E	I
DALLAS (DL-E)	RAVENIA (RV-E)	NY-:
3	E	T
FLORENCE (FL-E)	WASHINGTON D.C. (WA-E)	NY-T
E	E our-c	F-CLASS
BROADWRY (BR-E)	BOSTON (BS-E)	

Fig. 1. Input imagery used in the character recognition study.

TM-CTASS

evaluate the performance of the five discrete-valued filters. This study consisted of three phases. The first phase evaluated the performance of a BPOF made from a single character and the use of a TR-TPAF to increase the out-of-class discrimination. The second phase had two parts: (1) the use of fSDF-BPOF and pSDF-BPOF formulations to create filters designed to give equal-peak correlation response using only in-class training, and (2) the use of a TR-TPAF and pSDF-BPOF to maintain equal-peak correlation response within the in-class training set while discriminating against out-of-class characters, using both inclass and out-of-class training. In the third study phase, the TR-TPAF and the pSDF-BPOF formulations were used with a different training set to create filters with the same goal as the previous filters.

4.1. Character imagery and computer software

The characters used in these investigations were taken from an Apple Macintosh computer. Several different fonts of the character E and other characters from the New York font were used. Figure 1 displays all of the characters used in these investigations. The characters were 24 point bold and were generally 18 pixels high.

A 64-by-64 sample fast Fourier transform (FFT) and other operations were performed to simulate optical correlation and to analyze and display the correlation results.




Fig. 2. Discrimination of the SPOF for (a) the NY-E and (b) the Pi-E.

Fig. 3. Discrimination of the TR-TPAF for (a) the NY-E and (b) the

4.2. Definitions of filter performance metrics

Discrimination: The discrimination of a character by a given filter is the ratio of the maximum correlation peak intensity taken over all in-class inputs to the highest correlation peak intensity associated with the character.

Recognition: A character is recognized when its discrimination is less than 3 dB.

Correct in-class response is when an in-class character is recognized.

False in-class response (or false alarm) is when an out-ofclass character is recognized.

Correct out-of-class response is when an out-of-class character is not recognized.

Horner efficiency: In this study Horner efficiency is defined as the fraction of input scene energy appearing in a 5-by-5 pixel element block centered on the maximum correlation-peak sample. This block size was arbitrarily chosen based on the observation that correlation peaks typically were several pixels wide.

4.3. Character recognition studies

Two BPOFs were created from the Fourier transform of the New York E (NY-E) and Pierce E (PI-E), which were chosen arbitrarily. The performance of both filters was tested using the input images of the character E from 18 different fonts (in-class) and eight other characters from the New York font (out-of-class). The results of this study are shown in Fig. 2. As shown in Fig. 2(a), the NY-E BPOF recognized two inclass fonts, the NY-E and the Manhattan E (MN-E), and a out-of-class character, the NY-F. The BPOF made from the PI-E, as shown in Fig. 2(b), discriminated against all of the out-of-class characters and recognized only the PI-E character from the in-class character set.

To improve the (already good) discrimination performance of the BPOFs, ternary filters were created. A transform-ratio binary-amplitude mask was created using an in-class training set consisting of the NY-E, PI-E, MN-E, and the SG-E (Saigon) and an out-of-class training set consisting of the NY-F, -P, -R, and -B. Figure 3 shows the per-

TABLE I. Correlation performance of the tested filters using two-character input scenes.

Study	Out-of-class training set	Filter type	% pixels zeroed	Horner efficiency	Discrimination (dB)
1	N	NY-E BPOF	NA	8.30	1.81
1	N	PI-E BPOF	NA	7.10	4.80
1	Y	NY-E TPAF	90.28	0.03	5.34
1	Y	PI-E TPAF	90.28	0.10	> 10.00
2	N	ISDF-BPOF	NA	10.79	1.39
2	N	pSDF-BPOF	NA	8.40	1.35
2	Y	ITPAF	71. 36	0.27	3.14
2	Y	pSDF-BPOF	NA	6.81	1.73
3	Y	fTPAF .	84.23	0.18	3.21
3	Y	ITPAF	88.48	0.07	4 24
3	Y	pSDF-8POF	NA	2.52	0.85

TABLE II. Classification (discrimination) performance of the tested filters.

Study	Filter type	Out-of-class training set	% pixels zeroed	% correct in-class per in-class inputs	% false alarms per out-of-class inputs	% correct in-class per training set
1	NY-E BPOF	N	NA	11.11	12.50	NA
1	PI-E BPOF	N	NA	5.56	0.00	NA
1	NY-E TPAF	Y	90.28	50.00	0.00	NA
1	PI-E TPAF	Y	90.28	5.55	0.00	NA
2	ISDF-BPOF	N	NA	100.00	62.50	100.00
2	pSDF-BPOF	N	NA	72.22	62.50	58.33
2	TPAF	Y	71.38	100.00	0.00	100.00
2	pSDF-BPOF	Y	NA	77. 77	50.00	86.67
3	ITPAF	Y	84.23	60.00	0.00	100.00
3	TPAF	Y	88.48	80.00	0.00	100.00
3	pSDF-8POF	Y	NA	28.67	12.50	50.00




Fig. 4. Correlations using the NY-E filters of types (a) BPOF and (b) TR-TPAF. The input scene included the NY-E and NY-F. The far left peak is the NY-E subcorrelation.

formance of the TR-TPAFs for the NY-E and the PI-E. As the plots indicate, the application of the binary-amplitude mask to the BPOFs significantly improved the discrimination against the NY-F character by more than 3 dB. The NY-E TR-TPAF, although not designed to do so, increased the number of correct in-class responses from two to nine.

The BPOFs and the TR-TPAFs were tested with two input scenes, both containing one in-class character placed in the first quadrant and the out-of-class character, NY-F, placed in the fourth quadrant. The in-class character for the

first input scene was the NY-E, while for the second input scene it was the PI-E. The correlation response of the BPOF and TR-TPAF for the NY-E with the first input scene is shown in Fig. 4. The discrimination for the BPOF was only 1.81 dB, while the TR-TPAF produced discrimination above 5 dB. The correlation plane for the TR-TPAF exhibits increased noise around the correlation peak response compared with that for the BPOF. However, this should not be a problem for character recognition because an area (e.g., the size of the character) may be defined around the center

location of the character in which the maximum correlation response must fall to be considered a correct in-class correlation response.

Another important metric is the Horner efficiency. To create the TR-TPAFs, the transform-ratio method zeroed 90.28% of the filter elements. The Horner efficiency for the correlation response shown in Fig. 4(a) is 8.3%, while for the ternary filter it is only 0.034%. Since the BPOFs pass all of the energy of the input transform, they generally result in much better efficiencies than those filters that include amplitude modulation.


A summary of the performance of the PI-E BPOF and TR-TPAF for the second input scene is provided in the first section of Table I. As the table indicates, the discrimination performance of the TR-TPAF was greatly improved over that of the BPOF, but again the ternary filter produced very low Horner efficiencies.

The second and third phases evaluated the performance of the discrete-valued composite smart filters using two different training sets. The results of these investigations as well as the previous investigation are summarized in Tables I and II. Table I gives the correlation results using the two input scenes described earlier.

The in-class training set used in the second phase consisted of 12 characters, the NY-E, DL-E, FL-E, GE-E, CD-E, SF-E, PI-E, ST-E, MP-E, BR-E, TF-E, and CP-E, displayed in Fig. 1. The first part of this investigation used the fSDF-BPOF and the pSDF-BPOF formulations to create a filter that would maintain an equal-peak correlation response over the in-class training set. As mentioned, the creation of the fSDF-BPOF required the use of an iterative technique to solve for the weighting coefficients. The iterative technique suggested by Jared and Ennis15 was modified for our studies. Originally, this technique was formulated using only the zero-shift correlation response, but for this investigation the maximum response in the entire correlation plane was used. The reason is that the maximum cross-correlation response for the various members of the training set was not always the zero-shift response but would likely be detected in a practical system. The same issue was addressed by the correlation-SDF formulation.11

Table II lists the results of this investigation. As indicated by the table, the performance of the ISDF-BPOF met the specified filter formulation goal (recognition of all in-class inputs), while the pSDF-BPOF failed to do so.-Both the ISDF-BPOF and the pSDF-BPOF (incorrectly) recognized out-of-class characters.

A TR-fTPAF and pSDF-BPOF were created with out-of-class discrimination addressed in their formulations. The out-of-class training set included all eight of the out-of-class characters shown in Fig. 1. Both filters were tested with the characters displayed in Fig. 1. Figure 5 shows the plots of the correlation response for both filters. The results are also summarized in Table II. The TR-fTPAF met the design goals. It recognized all of the in-class characters and did not recognize any of the out-of-class characters. On the other hand, the pSDF-BPOF did not meet the goals of the filter formulation, recognizing only 66.7% of the in-class characters and properly classifying only 50% of the out-of-class characters. These filters were tested with one of the input scenes described earlier, containing the in-class character NY-E and the out-of-class character NY-F. The






Fig. 5. Discrimination of (a) the TR-ITPAF and (b) the pSDF-BPOF for the second phase of the study.

results are presented in Table I. Again, it is seen that there is a sacrifice in Horner efficiency for the improved discrimination performance of the TR-ITPAF.

In the third and final phase of our investigation, two TRfTPAFs and a pSDF-BPOF were created. Again, the goal of these filters was to produce an equal-peak correlation response for all members of the in-class training set while discriminating against all out-of-class characters. The inclass training set used in the formulation of these filters contained the NY-E, PI-E, WA-E, BS-E, MN-E, SD-E, RV-E, and SG-E.

The difference between the two TR-ITPAFs created was in the percentage of filter elements zeroed. One filter used a transform-ratio that zeroed 84.23% of the filter elements, while the other had 88.48% of the filter elements zeroed. The out-of-class training set used to create these filters included NY-F, -R, and -B. The pSDF-BPOF was created using an out-of-class training set that included the NY-F, -R, -B, -8, and -P.

The filters were tested with 15 in-class characters and all eight of the out-of-class characters. Figures 6 and 7 show the




Fig. 6. Discrimination for the TR-ITPAF with (a) 84.23% of the elements zeroed and (b) 88.46% zeroed.




Fig. 7. Discrimination of the pSDF-SPOF in the third study.




Fig. 8. Correlation for the TR-(TPAF having (a) 84.23% of the elements zeroed and (b) 82.48% zeroed, and (c) for the pSDF-8POF. The input scene included the NY-E and NY-F. The peak located at the far left corresponds to the NY-E input location.

plots of the correlation response of the filters. A summary of the results is provided in Table II. The TR-TPAFs performed quite well. Both met the specified filter formulation goal, while the pSDF-BPOF did not. These filters were also tested with the two-character input scene containing the NY-E and NY-F. Plots of the correlation output response are shown in Fig. 8. The output correlation plane is noisy, but again, for character recognition this may not be of great concern. The TR-TPAFs produced more than 3 dB of discrimination, but their Horner efficiency values were very low.

5. CONCLUSION

The goal of this study was to evaluate, by simulation, the performance of smart BPOFs and ternary filters for in-class recognition and out-of-class discrimination. The performance of these filters for specific character recognition problems was addressed. As Tables I and II indicate, the ter-

jary filters for the cases studied were capable of achieving the highest discrimination. The TR-ITPAFs were the only filters capable of meeting the in-class character recognition and out-of-class discrimination goals set for these studies. As indicated in Table I, the improved performance of the ternary filters comes with a significant reduction in Horner efficiency, and this may be viewed as the price of this performance. The practical implications of these lower efficiencies must be evaluated experimentally.

6. ACKNOWLEDGMENT

This work was supported by the U.S. Air Force Rome Air Development Center, Hanscom AFB, Mass., under contract F19628-87-C-0073.

1. REFERENCES

- 1. D. Psaltis, E. G. Pack, and S. S. Venkatesh, "Optical image correla-tion with a binary spatial light modulator," Opt. Eng. 23(6), 698-704 (1964).
- W. E. Ross, D. Psaltis, and R. H. Anderson, "Two-dimensional magneto-optic spatial light modulator for signal processing," in Real Time Signal Processing V, J. Trimble, ed., Proc. SPIE 341, 191-196 (1982).
- J. L. Horner and H. O. Bartelt, "Two-bit correlation," Appl. Opt. 24(18), 2889-2893 (1985).
- D. L. Flannery, J. S. Loomis, M. E. Milkovich, and P. E. Keller, "Application of binary phase-only correlation to machine vision," Opt. Eng. 27(4), 309-320 (1965).

 J. Horner, "Light utilization in optical correlators," Appl. Opt.
- 21(24), 4511-4514 (1982).
- P. D. Gianino and J. L. Horner, "Additional properties of the phase-only correlation filter," Opt. Eng. 23(6), 693-697 (1984).
 G. Schils and D. Sweeney, "Iterative technique for the synthesis of distortion-invariant optical correlation filters," Opt. Lett. 12, 307
- A. Mahalanobis. B. V. K. Vijnya Kumar, and D. Cassesst, "Spatial-temporal correlation filter for in-plane distortion invariance," Appl. Opt. 25(23), 4466-4472 (1986).
- "Unified synthetic discriminant function computational D. Cassent.
- D. Casseett, "Unified synthetic discriminant runction computational formulation," Appl. Opt. 23(10), 1620–1627 (1984).

 D. Cassett, W. Rozzi, and D. Fetterly, "Projection synthetic discriminant function performance," Opt. Eng. 23(6), 716–720 (1964).
- (1903).

 Cassesent and W. Chang, "Correlation synthetic discriminant functions," Appl. Opt. 25, 2343 (1996).

 B. V. K. Vijaya Kumar, E. Pochapsky, and D. Cassesent, "Optimality considerations in modified matched spatial filters," in Analog Optical Processing and Computing, H. J. Caulfield, ed., Proc. SPIE 519, 85-93 (1985)
- 13. J. L. Horner and P. D. Gianino, "Applying the phase-only filter concept to the synthetic discriminant function correlation filter," Appl.
- Opt. 24, 851 (1985).
 R. Kaliman, "Direct construction of phase-only correlation filters,"
- Appl. Opt. 26(24), 5200 (1967). D. Jared and D. Ennis, "Inches ion of filter modulation in synthetic-
- D. Jared and D. Ennis, "Lacusates of their mountains in synthetic-discriminant function filters," Appl. Opt. 28, 232 (1989).
 D. A. Jared, D. J. Ennis, and S. A. Drestin, "Evaluation of binary-phase-only-filters for discertion-invariant pattern recognition," in Computer-Generated Holography II, S. H. Lee, ed., Proc. SPIE 884, 139-145 (1988).
- D. L. Flannery, J. S. Losenia, and M. E. Milkovich, "Transformatio ternary phase-amplitude filter formulation for improved correlation discrimination," Appl. Opt. 27(18), 4079-4083 (1988).

- N. A. Clark and M. A. Handschy, "Device application of ferroelectric liquid crystals," in Optical Computing, J. A. Nell, ed., Proc. SPIE 625, 60-61 (1986).
- R. R. Kallman, "Optimal low noise phase-only and binary phase-only optical correlation filters for threshold detectors," Appl. Opt. 25(23),
- 20. D. L. Flannery, J. S. Loomis, and M. E. Milkovich, "Design elements of binary phase-only correlation filters," Appl. Opt. 27(20), 4231-4235 (1988).
- 4231-4235 (1988).
 D. Cottrell, R. Lilly, J. Davis, and T. Day, "Optical correlator performance of binary phase-only filters using Fourier and Hartley transforms," Appl. Opt. 26, 3755 (1987).
 D. Flannery, J. Loomis, and M. Milkovich, "New formulations for discrete-valued correlation filters," in Digital and Optical Shape Representation, R. D. Juday, ed., Proc. SPIE 938, 206-211 (1988).



Mary E. Milkovich received her BS degree in physics in 1986 from John Carroll University and an MS degree in electro-optical engineering in 1968 from the University of Dayton. She is currently a developmental engineer in the Federal Systems Division of Eastman Kodak Co. in Rochester, N.Y. She is a member of Sigma XI and SPIE.



David L. Flannery received the Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology in 1968 and is a senior research engineer in the Applied Physics Division of the University of Dayton Research Institute. He has 20 years' experience performing and managing research in the areas of optical processing and computing and spatial light modulators.



John S. Loomie has been a research optical physicist and professor of electro-optics at the University of Dayton since 1979. His research interests include image analysis, computer graphics, geometric optics, interferometry, and ellipsometry. He has authored computer programs in interferogram reduction, optical design, and thin-film design and analysis. Dr. Loomis received his BS degree in physics from Case institute of Technology, an MS degree in physics from the University of Illinois, and his

Ph.D. degree in optical sciences from the University of Arizona. He is a member of the Machine Vision Association of the Society of Manufacturing Engineers, the Association for Computing Machinery. OSA, and SPIE

- G Neugebauer, R. Hauck, and O. Bryngdahl, "Computer-Generated Holograms: Carrier of Polar Geometry," Appl. Opt. 24, 777 (1985)
- H. Onda, "High Precision Positioning System Using Rectangular Thrust Air Bearing," in Proceedings, IECON'85, International Conference on Industrial Electronics, Control and Instrumentation, San Francisco, CA (18-22 Nov. 1985), p. 321.

Implementation of ternary phase amplitude filters using a magnetooptic spatial light modulator

Brian A. Kast, Michael K. Giles, Scott D. Lindell, and David L. Flannery

Brian Kast and Michael Giles are with New Mexico State University, Department of Electrical & Computing Engineering, Las Cruces, New Mexico 88003; the other authors are with University of Dayton Research Institute, Applied Physics Division, Dayton, Ohio 45469.

Received 5 December 1988.

Sponsored by Thomas K. Gaylord, Georgia Institute of Technology.

0003-6935/89/061044-03\$02.00/0.

© 1989 Optical Society of America.

Ternary modulation can be achieved with a single SLM and improved correlation discrimination can be experimentally realized.

The use of a ternary phase amplitude filter (TPAF) to improve the performance of binary phase-only correlators has been proposed, analyzed, and simulated. This Communication describes the experimental implementation of a simple TPAF using a 128 × 128 element SIGHT-MOD (commercially available through Semetex Corp.) magnetooptic spatial light modulator (MOSLM). A brief description of the experiments, which included spatially resolved transmission measurements and TPAF correlations, and a comparison of experimental and theoretical (computer simulation) correlation results are presented.

The ideal TPAF superimposes a zero modulation state on a binary phase-only filter, the encoding the modulations -1, 0, and 1. The zero modulation state is used to block spatial frequency components and otherwise would degrade

correlation performance.

The MOSLM is generally used as a binary device. Its operation in this mode has been reported in detail.^{2.3} To review, a normal write sequence involves three steps—erase, nucleate, and saturate. First, all pixels are set to one of two stable saturated magnetic states (erased) by applying a bulk magnetic field generated by an external coil. Second, the magnetization state of selected pixels is partially reversed (nucleated) by passing current pulses through the appropriate conducting traces on the device. Finally, the nucleated pixels are driven to complete saturation in the reversed magnetic state (written) by a second application of the coilgenerated magnetic field (in the opposite direction from the erase coil pulse). A following polarizer oriented perpendicular to the linear polarization of an incident beam will resolve the oppositely (Faraday) rotated fields of the transmitted light emerging from oppositely programmed pixels into components which differ by 180° phase, thus effecting -1 and 1 modulation states.

A third stable pixel magnetization state, comprising mixed magnetization, was reported² and is the basis for achieving the zero modulation state. In the third state, different regions of a pixel are randomly distributed between the two

magnetization directions, with approximately equal probability. This state is achieved by operating the device exactly as described except that an additional nucleation step is applied to selected erased pixels not previously written. The zero modulation state is simply the result of nucleating a pixel without applying a subsequent coil-generated field to

carry out saturation.

The characteristic magnetic domain size in the mixed magnetization state averages less than half of the dimension of the active pixel area.² The effective modulation values within each domain are still -1 and +1 and thus average intensity transmission of a mixed pixel should equal that of a saturated pixel. However the small characteristic dimension is expected to result in most of the transmitted light being diffracted outside the Nyquist spatial frequency bandpass normally used by a BPOF correlator (defined by the pixel spacing), thus effectively comprising a zero modulation state.

To investigate the ternary state characteristics of the MOSLM, 128×128 devices were incorporated at the input plane in coherent imaging systems with bandpass apertures at the Fourier planes sized to pass spatial frequencies only up to half of the Nyquist frequency, corresponding to the maximum alias-free frequency that can be written on the devices. In a first experiment the large area transmission of the three modulation states was measured by collecting light imaged through a 20×20 pixel region in the center of the MOSLM with a photodiode. When the polarizer was set to achieve equal intensity transmission with all elements programmed to -1 or +1 modulation, transmission was reduced to 8% of the transmission of either phase modulation state with all elements in the zero modulation state.

When the bandlimiting aperture was removed, transmission of the third state increased to 45%, consistent with the hypothesis that high-order diffraction is the mechanism for zero modulation.

Spatially resolved transmission measurements were performed with the bandlimiting aperture by reading out the image of the MOSLM with an RCA TC1030 SIT (silicon intensified target) video camera (without lens), digitizing the image, and plotting image slices with offset and gamma correction. The MOSLM was programmed with a series of vertical bars representing repeating cycles of -1, 0, and 1 modulation. Traversing the MOSLM, the first cycle used 8 pixels for each bar, and each successive cycle used one less pixel per bar, down to one pixel per bar, which was repeated seven times. Figure 1 is a video image of this pattern.

Figure 2(a) is a slice plot across the image portion containing the larger bars. Prior careful measurements indicated a gamma of ~0.7 for the camera, in agreement with manufacturer's specifications, and these results were corrected for that value. Offset correction was referenced to black areas surrounding the MOSLM image and by using opaque bars to create shadows across the image. Black level was found to be constant over the image so offset correction was simplified. The plots were averaged over 15 adjacent horizontal scan lines to achieve noticeable, but not major, noise reduction.

Excellent contrast between unity and zero modulation states is apparent. Zero modulation state transmission is estimated to be <2.6% of unity magnitude state transmissions based on averaging over more than ten samples in the appropriate regions. The dips between -1 and +1 modulation bars are believed consistent with the amplitude zero crossings which occur at those points. Significant coherent imaging noise, and probably some electronic noise, are believed to contribute to data fluctuations seen in the plots, although a detailed theory might predict some of the artifacts.





Fig. 1. Video image of MOSLM programmed with ternary bars (+1, 0, -1) of various widths.






Fig. 2. Slice plots across the pattern in Fig. 1. (a) Coarse bars (eight pixels/bar at left down to two pixels/bar at right). 'b) Fine bars (four pixels/bar at left down to one pixel/bar at right). The sloping envelope of peak values in left portion of (a) is due to nonuniform illumination of the MOSLM.

Figure 2(b) is a slice through the region of one pixel per har cycles, with an expanded scale relative to Fig. 2(a). Only the fundamental frequency (corresponding to a period of three MOSLM pixel spacings) of this pattern is passed by the bandlimiting aperture. If the MOSLM zero state is not perfect, or the polarizer is not set for balanced unity modulation states, some dc component is also passed. The pure fundamental, with no dc component, yields a sine wave of twice the original spatial frequency, ranging in intensity from zero to a maximum, after square-law detection. With some dc component, the detected intensity would exhibit an increase (decrease) of every other peak. The spatial frequency of the detected fundamental is believed sufficiently high that its detection by the camera is significantly degraded by the video system resolution. Simple Fourier series analysis indicates the average value of the transmission in this region should be 34% of the unity transmission average for the larger bars and, with no resolution or sampling limitations, the wave should oscillate from zero to twice the average intensity. The effect of limited video resolution is to reduce the amplitude of the oscillating component but not the average level.

We believe the results shown are consistent with the picture of a detected fundamental sine wave degraded by system resolution. Alternating peaks are approximately equal in intensity and the average is about the level predicted by theory. However the data of Figs. 2(a) and (b) were taken with slightly different polarizer settings. The setting yielding equal alternate peaks of the detected one bar per pixel fundamental differed from that yielding equal intensities of (larger) +1 and -1 bars by $\sim 0.5^{\circ}$. We hypothesize that this is a manifestation of a slight dc component due to a nonideal zero modulation state which is likely since there is no guarantee the mixed state domain areas are exactly equal in a pixel.

Experimental correlations were performed in a converging beam configuration with 128 × 128 MOSLMs used both for input and as the spatial filter. These experiments were performed to demonstrate that a simple TPAF can improve performance of a binary phase-only correlator. The input MOSLM was operated in the binary amplitude mode, and the filter MOSLM implemented BPOFs or TPAFs.

A binarized image of an F16 aircraft degraded by computer-generated clutter was used as the input scene (Fig. 3). A Hartley BPOF^{4,5} was computed from the uncluttered target image. A TPAF was generated by setting a central 5×5 square of the BPOF elements to zero. These pixels attenuate the lowest spatial frequency components of the input image, thus removing some of the degrading effects of the clutter.

Computer simulations of the correlations were performed using TPAF modulation values of -1, 0.1, and 1. Results of experimental correlations and computer simulations are pre-



Fig. 3. (a) Cluttered input image and (b) object from which the filters were generated.




Fig. 4. Profiles of correlation peaks: (a) optical BPOF, (b) optical low frequency block TPAF, (c) and (d) computer simulations of (a) and (b).

sented in Fig. 4. Signal-to-noise ratios (SNR) measured for the 2-D correlation patterns whose profiles are shown in Fig. 4 indicate the TPAF provides a SNR improvement of 1.6 dB for the simulation and 1.55 dB for the experiment. SNR is defined as the ratio of peak correlation value to the rms value of all points in the correlation outside the peak's 3-dB width.

These results verify that useful ternary (-1, 0, 1) modulation can be achieved with a single SLM (the MOSLM) and that the improved correlation discrimination theoretically predicted for TPAFs can be experimentally realized.

References

- D. L. Flannery, J. S. Loomis, and M. E. Milkovich, "Transform-Ratio Ternary Phase-Amplitude Filter Formulation for Improved Correlation Discrimination," Appl. Opt. 27, 4079 (1988).
- W. E. Ross, D. Psaltis, and R. H. Anderson, "Two-Dimensional Magneto-Optic Spatial Light Modulators for Signal Processing," Proc. Soc. Photo-Opt. Instrum. Eng. 341, 191 (1982).
- 3. D. Psaltis, E. G. Paek, and S. S. Venkatesh, "Optical Image Correlation with a Binary Spatial Light Modulator," Opt. Eng. 23, 698 (1984).
- R. V. L. Hartley, "A More Symmetrical Fourier Analysis Applied to Transmission Problems," Proc. IRE 30, 144 (1942).
- D. M. Cottrell, R. A. Lilly, J. A. Davis, and T. Day, "Optical Correlator Performance of Binary Phase-Only Filters Using Fourier and Hartley Transforms," Appl. Opt. 26, 3755 (1987).

This work was supported by Sandia National Laboratories contract 06-0436, by U.S. Army Research Office contract DAAL03-87-K-0106, and by U.S. Air Force Rome Air Development Center contract F19628-87-C-0073.

The authors wish to thank William Ross of Litton Data Systems, Van Nuys, CA, and Jaye Waas of Semetex Corp., Torrance, CA, for helpful discussions.

A REMINDER

Authors submitting Rapid Communications for publication must remember to include a letter from their institution agreeing to honor the publication charge. Otherwise publication will be delayed and the whole idea of this section of Latters to the Editor is defeated.

APPENDIX G

Ternary Phase-amplitude Filters for Character Recognition

Scott D. Lindell* and David L. Flannery The University of Dayton Research Institute Dayton, Ohio 45469

ABSTRACT

Ternary phase-amplitude filters (TPAF, encoding the modulation states -1, 0, and 1) are being developed for correlation both in theory and in practice. They are attractive due to their few discrete modulation levels which facilitate efficient electronic storage for on-line systems and allow implementation in real time with available spatial light modulators (SLMs), as recently demonstrated using magneto-optic devices. Simulations have demonstrated that effective smart TPAFs can be formulated to address class discrimination goals. We report experimental demonstration of significant increases in correlation discrimination achieved with TPAFs designed with the transform ratio technique. Results are in substantial agreement with theory based on computer simulations, and verify the practical implementation of improved discrete-level filters using magneto-optic SLMs.

1. INTRODUCTION

Matched spatial filtering is a method used for pattern recognition applications such as character and target recognition. Implementation of the classical matched filter in a Fourier optical processor requires simultaneous and independent modulation of the phase and magnitude of an optical Fourier transform. This requirement and the intrinsic distortion sensitivity of the matched filter have limited its practical (real-time) application. However, binary phase-only versions of the matched filter (BPOF) have been implemented successfully in computer simulations [1,2,3] and in real-time optical correlators using commercially available spatial light modulators (SLMs) [4,5]. The ternary phase-amplitude filter (TPAF) [6,7] is a discrete-valued formulation that includes simultaneous and independent binary magnitude and binary phase modulation. The TPAF has been shown, in computer simulations, to provide improved discrimination over that of the BPOF. "Smart" (distortion-invariant) TPAF formulations also have demonstrated good performance in simulations [7,8]. Recently it was demonstrated that the magneto-optic spatial light modulator (MOSLM) developed and manufactured by Litton and Semetex Corporations is capable of ternary modulation suitable for TPAFs [9]. This work used the sero-state of the filter only to implement a simple DC block. The resulting improvement in correlation signal-to-noise, although significant, is small compared to what should be possible with more complicated sero-state patterns, as indicated by reported simulations [6,7,8].

The work reported here, involving both computer simulations and matching experimental results, demonstrates that the TPAF formulation provides improved discrimination over the BPOF for both character and target recognition, using zero patterns computed by the transform-ratio technique [6]. Practical limitations, including those due to a nonideal zero-magnitude SLM state, were investigated and will be discussed.

2. BPOF AND TPAF FILTERS

The TPAF evolved from the binary phase-only filter (BPOF), which is of interest for use in optical correlation systems because it can be effectively implemented with commercially available MOSLMs and exhibits good correlation performance. The BPOF is a discrete-valued subset of the phase-only filter (POF) in which only two phase modulation values, 0 and π in phase, or 1 and -1 in amplitude, are encoded [1,2].

Horner and Leger [2] showed with computer simulations that objects could be Fourier transformed and binarised by thresholding on the imaginary component of the transform with the resulting BPOF capable of pattern *Currently employed by Martin Marietta Astronautics, Denver, Colorado 80201

recognition. Psaltis et al [1], thresholding on the real component of the Fourier transform, implemented a BPOF in an experimental optical correlator using the MOSLM as the Fourier plane modulator. These two choices of thresholding axis, or threshold line angle (TLA) [3], are the extremes of a continuum used to construct a BPOF as follows:

 $F_{\mathcal{B}}(u,v) = \begin{cases} +1 & \text{Re}[F(u,v)] \cos \theta - \text{Im}[F(u,v)] \sin \theta \ge 0 \\ -1 & \text{otherwise} \end{cases} , \tag{1}$

where θ defines threshold line angle in the complex plane of the Fourier transform.

Advantages of the BPOF over the classical matched filter include the easy and successful implementation of the BPOF in real-time, the minimisation of memory required to store a filter, and sharper correlation peaks facilitating target location estimation. Disadvantages of the BPOF (relative to a classical matched filter) include increased sensitivity to geometric distortions of the input, reduced signal-to-noise in many cases, and passing all input plane energy (including that of clutter and nontargets) to the correlation plane.

The ternary phase-amplitude filter (TPAF) [6] concept was developed to improve the discrimination capabilities over that of the BPOF or other smart filters using binary phase modulation by adding binary magnitude modulation. Since binary phase filters modulate only the phase of the input's Fourier transform, the energy of all spatial frequency components of the input is transmitted through the spatial filter. Hence the fundamental limit in reducing the cross correlation responses, caused by non-target inputs, is the even distribution of non-target energy across the correlation plane. The encoding of the TPAF modulation states (1, 0, -1) couples the BPOF's (1, -1) modulation states with binary magnitude (1, 0) modulation to provide a means of attenuating spectral regions of the input most closely related to non-target information, while retaining recognition via the phase states and the potential for practical real-time implementation with existing SLMs due to the discrete modulation levels.

TPAF modulation can be viewed as a multiplication of a BPOF with a binary magnitude mask (BMM). A DC block is an example of a simple BMM. In general, BMMs are non-trivial patterns formulated by a design procedure such as the transform ratio (TR) technique [6]. The TR technique uses representative power spectral density patterns of both the in-class and out-of-class training sets. The value for each resolvable frequency region (e.g., pixel) of the BMM is determined by thresholding on a ratio computed from the corresponding spectral energies representing the two training sets. In this study the computed ratio was simply the in-class power spectral density divided by that of the out-of-class pattern(s). The TR method can be described as follows: if the ratio of the power spectral density of the in-class training set, $|I(u,v)|^2$, to that of the out-of-class, $|O(u,v)|^2$, at a particular frequency region is greater than some threshold, T, the corresponding TPAF value, F(u,v), is set to the phase modulation state defined by the BPOF; otherwise the TPAF filter value is given a sero-amplitude modulation state. To show this let

$$R(u,v) = \frac{|I(u,v)|^2}{|O(u,v)|^2},$$
 (2)

then

$$F(u,v) = \begin{cases} \pm 1 \text{ (per BPOF)} & R(u,v) \ge T, \\ 0 & \text{otherwise} \end{cases}$$
 (3)

3. RESULTS

Implementing the TPAF in a real-time correlator requires an SLM capable of simultaneous and independent binary amplitude and binary phase modulation. This modulation has been achieved in the MOSLM with the sero-amplitude modulation state intensity transmission being approximately 3.0% in both 128x128 [9] and 48x48 devices. Implementation of a 5x5-pixel DC block in a 128x128 MOSLM correlator using the sero-amplitude modulation state has been demonstrated by Giles and coworkers with a SNR improvement of 1.55 dB over that obtained with a BPOF. [9]

To demonstrate the implementation of non-trivial TR TPAFs in a real-time 48x48 MOSLM correlator, three different inputs and five different references were considered; see Figure 1. Two of the inputs, the "X-O" and "E-F"




Figure 1: Binary input patterns used in TPAF correlation study.

patterns represent character recognition problems while the "F1-F2" input pattern represents target recognition of two different airplanes. All the images used were computer-generated and involve no background noise. In all cases reported here, inputs included two of the patterns while the reference was based on one of these two patterns. Thus the latter pattern was the in-class object while the other input pattern was the out-of-class object, and the primary emphasis was on improving discrimination between the two input objects using the TR technique to set the sero-state pattern of the filter.

The TR threshold design parameter should be optimised for each particular input/reference case and determines the percentage of filter elements modulated with the sero-amplitude state. Practical considerations such as the Horner efficiency [10,11], the input laser power, and the leakage of the sero-amplitude modulation state pixels require attention when the filter is being optimised for use in an experimental correlator.

Initial TPAF computer simulations were performed using a model that assumed an ideal (opaque) sero-amplitude modulation state. The TPAF and BPOF results of the experimental and simulated studies are shown in Table 1 which also includes other simulation results to be discussed later. The table values show that for all five cases studied the implementation of the TPAF provided significantly improved discrimination over that obtained with the BPOF. Particularly notable improvements occurred for the "F2" and the "E" references correlated with the "F1-F2" and "E-F" inputs respectively. Horner efficiency values from the simulations are given in Table 1, and are computed with a modified definition: the quotient of the peak correlation pixel energy and the total input energy of the target (reference) pattern. Use of the single peak pixel is justified on the basis that most of the correlation peak energy is concentrated in a single pixel. Use of only the input target energy as the base for efficiency is justified on the basis that it is inappropriate to penalise the rating of a filter because of the presence of clutter in the input scene, as would be done using the original definition of Horner efficiency [10]. Peak-to-clutter ratio was defined in terms of the highest correlation peak pixel intensity divided by that of the next highest peak in the output plane.

Using the "F2" as the in-class object, a TPAF with a blocking percentage of 58.4% provided an experimental peak-to-clutter (P/C) improvement from 2.34 dB for a cosine-BPOF to 5.13 dB. A typical criterion for acceptable discrimination by a filter is a P/C value of greater than 3.0 dB; thus the BPOF recognised both input objects in the scene as being the reference, "F2", while the TPAF rejected the out-of-class character. A TPAF using the "E" as the in-class object in the "E-F" input pattern with a blocking percentage of 67.0% yielded an improvement in P/C from 1.60 dB to 4.99 dB. Again the BPOF using the "E" as the reference recognised both objects of the input as being "E's" while the TPAF recognised only the proper in-class input.

The other three cases studied show similar experimental discrimination improvements with the TPAF formulations. Qualitative improvement of the TPAF over the BPOF can be seen by comparing (linear-in-intensity) plcts of the simulated and experimental correlation responses such as those shown in Figure 2. In this figure the target "E" response is located at the far rear while the clutter "F" response is located forward right. Note the substantial suppression of the clutter response in both theory and experiment with the TPAF.

A study of the relationship between peak-to-clutter and blocking percentage of TPAFs was conducted for three different references. The purpose of the study was to investigate the general agreement between simulation and experimentation for a wide range of blocking percentages. The blocking percentages were chosen in an attempt to sample a large portion of the percent-seroed range with a concentration of cases in high Horner efficiency $(\eta_{H.E.})$ (i.e. lower percent-seroed) regions to better accommodate experimental procedure. The results of the comparison between simulation and experiment have been summarised in Figures 3–5. Figure 5 shows an additional simulation plot which will be discussed later.


For all three studies the experimental P/C values follow the trends of the ideal simulation P/C quite well. The trends indicate that increasing the blocking percentage provides improvements in the discrimination yielded by the filters. However, the Horner efficiency of the filter decreases as the blocking percentage increases, placing a practical limit on the P/C improvements observed in the experimental system. This is due to the system's increased susceptibility to extraneous noise at lower Horner efficiencies. This is seen in the results at higher blocking percentages where there is less agreement between the ideal simulation and experiment. The most likely

Table 1: TPAF and BPOF: Experimental and Simulated Results

Reference	% Zeroed	Experimental	Ideal Model	7H.B.	Ternary
			Simulation	(Ideal)	Leak. Model
1		P/C (dB)	P/C (dB)	%	P/C (dB)
X		-/ ()	1/3/32/		
[BPOF	7.88	7.97	30.9	
[56.4	9.26	12.4	10.8	9.63
0	00.2	0.20		20.0	0.00
	BPOF	7.38	7.77	29.8	
	63.5	8.59	12.2	4.30	10.5
F2	00.0		3210		20.0
[BPOF	2.34	2.55	12.0	
[35.7	2.66	3.08	9.20	2.98
	46.0	4.03	4.42	5.34	4.33
	50.7	4.28	4.45	4.61	4.25
	58.4	5.13	6.86	2.64	6.17
	70.1		8.33	1.30	8.90
	75,7		8.60	0.85	8.62
	82.4		8.95	0.49	7.75
	89.1		10.2	0.19	6.17
	91.6		9.51	0.09	5.51
F1					
	BPOF	5.27	5.34	10.8	
	41.6	5.44	6.00	6.25	6.68
	54.8	6.20	6.73	3.66	8.60
	58.4	7.86	8.78	2.38	9.86
	64.3	8.39	9.90	1.35	10.4
	72.7		11.0	0.50	9.83
	80.7		10.1	0.25	9.06
	86.8		7.54	0.09	5.76
	89.7		6.50	0.06	5.86
E			ļ		
	BPOF	1.60	1.73	15.2	
	21.6	2.22	2.53	12.4	2.36
	31.7	2.95	2.94	9.19	2.72
	40.5	3.37	3.41	7.31	3.04
	58.3	4.40	5.09	2.72	4.49
	67.0	4.99	6.38	1.44	5.33
	83.2		9.06	0.31	6.87
	91.7		5.87	0.06	4.60



Experimentation






Figure 2: Simulation and experimental correlation intensity patterns. (a) BPOF, (b) TPAF with 67% of elements served




Figure 3: Ideal simulation and experimental curves of P/C vs % Zero for "F2" reference.




Figure 4: Ideal simulation and experimental curves of P/C vs % Zero for "F1" reference.




Figure 5: Ideal and mixed-state leakage simulations, and experimental curves of P/C vs % Zero for "E" reference.

cause for this increase in disagreement is that the simulation model does not include the known leakage by the sero-amplitude modulation pixels in the MOSLM. This effect will become more prominent at high blocking percentages as a larger percentage of the total light reaching the correlation plane will not be predicted by the ideal model.

The leakage of light through a blocking pixel is due to unequal areas of the +1 and -1 states inside the pixel. This effect causes a DC component proportional to the difference in their areas to leak to the correlation plane where it is not attenuated by the Nyquist aperture comprised by the correlation plane. This DC component has a polarity identical to the state of larger area in the mixed-state pixel. For the case of the Fourier plane MOSLM used in this study, the leakage amplitude was approximately 0.17 of the transmittance of ± 1 pixels, with a polarity identical to that of the fully saturated +1 pixels. (Designation of pixel polarity is an arbitrary choice; it is only important to be consistent.) This leakage value implies that the area distribution in a mixed-state pixel is 53.0% for the +1 state and 47.0% for the -1 state in the device used.

The results of including the sero-state leakage in the computer model can be seen in Table 1. Table 1 indicates that the model including the leakage phenomena provides improved agreement with the experiment for four out of the five cases studied. The case that does not show improvement also has the lowest Horner efficiency, $\eta_{H,E}$ and thus may be affected more by phenomena not included in the model. These effects may include multiple reflections between components lacking AR coatings, other nonideal MOSLM characteristics, and general system noise. The effect of including the leakage in the simulation model is demonstrated graphically in Figure 5 for the case of the "E" as the reference for the "E-F" input pattern.

4. SUMMARY

BPOF and TPAF formulations were investigated both in computer simulations and in an optical correlator for several different in-class and out-of-class objects representing both character and target recognition. Experimental P/C improvements of as much as 3.6 dB, over corresponding BPOF's, were obtained for TPAFs with blocking percentages as high as 67.0%. The degree of agreement between simulation and experiment was generally good

but varied from case to case. Trends showed that at low blocking percentages the agreement was very good while agreement weakened at high blocking percentages. Inclusion of the leakage effect in the simulation model yielded improved agreement for all but one case studied.

From this study it is evident that the TPAF formulation provides an impressive improvement in discrimination over the corresponding BPOF while requiring no additional components in a MOSLM correlator. Implementation using the MOSLM involves only straightforward modifications of drive software, thus enabling easy real-time implementation. The agreement between experiment and simulation using a model incorporating the zero-amplitude modulation leakage is good, indicating that useful predictions of filter performance may be obtained by computer simulation.

5. ACKNOWLEDGMENT

This work was performed under Air Force contract F19628-87-C-0073 sponsored by the Rome Air Development Center, RADC/ESOP, Hanscom AFB, Mass.

6. REFERENCES

- [1] D. Psaltis, E. Paek, S. Venkatesh, "Optical Image Correlation with Binary Spatial Light Modulator," Optical Engineering, pp. 698-704, Nov./Dec. 1984.
- [2] J.L. Horner, J.R. Leger, "Pattern Recognition with Binary Phase-Only Filters," Applied Optics, pp. 609-611, Vol. 24, 1 March 1985.
- [3] D. Flannery, J. Loomis, and M. Milkovich, "Design elements of binary phase-only correlation filters," Applied Optics, Vol. 27, pp. 4231-4235, 15 October 1988.
- [4] D. Flannery, J. Loomis, M. Milkovich, and P. Keller, "Application of binary phase-only correlation to machine vision," Optical Engineering, Vol. 27, pp. 309-320, April, 1988.
- [5] S. Mills and W. Ross, "Dynamic magneto-optic correlator: real-time operation," Proc. SPIE Vol. 753, pp. 54-63, January 1987.
- [6] D. Flannery, J. Loomis, and M. Milkovich, "Transform-ratio ternary phase-amplitude filter formulation for improved correlation discrimination," Applied Optics, Vol. 27, pp. 4079-4083, 1 October 1988.
- [7] M. Milkovich, D. Flannery, and J. Loomis, "A Study of Transform-ratio Ternary Phase-Amplitude Filter Formulations for Character Recognition," Optical Engineering, Vol. 28, pp. 487-493, May 1989.
- [8] D. Flannery, J. Loomis, and M. Milkovich, "New formulations for discrete-valued correlation filters," Proc. SPIE, Vol. 938, pp. 206-211, April, 1988.
- [9] B. Kast, M. Giles, S. Lindell, and D. Flannery, "Implementation of Ternary Phase-amplitude Filters Using a Magneto-optic Spatial Light Modulator," Applied Optics, Vol. 28, pp. 1044-1046, 15 March 1989.
- [10] J.L. Horner, "Light Utilisation in Optical Correlators," Applied Optics Vol. 21, pp. 4511-4514, 15 December 1982.
- [11] J. Horner and P. Gianino, "Applying the phase-only filter concept to the synthetic discriminant function correlation filter," Applied Optics, Vol. 24, pp. 851-855, 15 March 1985.

Experimental investigation of transform ratio ternary phaseamplitude filters for improved correlation discrimination

Scott D. Lindell, MEMBER SPIE Martin Marietta Strategic Systems Denver, Colorado 80201

David L. Flannery University of Dayton Research Institute Dayton, Ohio 45469

Abstract. Transform ratio ternary phase-amplitude filters (TR-TPAFs) encoding the modulation states 1, 0, -1 have been investigated in theory and practice. Simulations have demonstrated that TR-TPAFs can be formulated to provide increased discrimination compared with binary phase-only filters. TPAFs have practical advantages over complex-valued matched filters-efficient electronic filter storage and real-time implementation with available devices, such as magneto-optic spatial light modulators (MOSLMs). These are associated with the use of only three discrete modulation levels. In an experimental study, significant increases in correlation discrimination were demonstrated for several test patterns, showing substantial agreement with computer simulations. The effects of an imperfect MOSLM zero-modulation state, present in our experiment, were modeled and investigated. The performance of the TR-TPAFs for different threshold line angles also was investigated.

Subject terms: optical pattern recognition; optical signal processing; correlation. ternary phase-amplitude filters.

Optical Engineering 29(9), 1044-1051 (September 1990)

1. INTRODUCTION

CONTENTS

- 1 Introduction
- 2. Background
- 3. Optical correlator system
- 4 Results
 - 4.1. Results compared to simulations using ideal zero state in the filter
 - 4.2. Modeling zero-state leakage
 - 4.3. Effects of inverting filter phase states relative to leakage
- 4.4. Threshold line angle effects
- 5. Summary and conclusions
- 6 Acknowledgments
- 7. References

Invited Paper PR-120 received March 7, 1990; revised manuscript received May Invited Paper PK-120 received March /, 1990; revised manuscript received may 1, 1990, accepted for publication June 21, 1990. This paper is a revision of paper 1151-22; presented at the SPIE conference Optical Information Processing Systems and Architectures, Aug. 8–11, 1989, San Diego, Calif. The paper presented there appears (unreferred) in SPIE Proceedings Vol. 1151. © 1990 Society of Photo-Optical Instrumentally a Engineers.

advanced TPAF formulations 5.6 have demonstrated significant

The evolution of real-time imaging sensor capabilities has cre-

^{*}Author was with the University of Dayton Research Institute when this work was performed.

ated a need for an advanced image processing technique for recognizing and locating patterns of interest. Coherent optical correlation is an analog image processing technique that has long been recognized as having the potential for implementing pattern recognition functions. The advancement of spanal light modulator (SLM) technology and a better understanding of spatial filters encoding limited modulation levels have been combined to create a real-time optical pattern recognition system concept capable of practical near-term implementation. The binary phase-only filter (BPOF) is a discrete-valued spatial filter that has been successfully implemented in real-time optical correlators using commercially available SLMs, exhibiting significant agreement with corresponding computer simulations. The ternary phase-amplitude filter (TPAF) also is a discrete-valued formulation. It includes simultaneous and independent encoding of binary phase and binary amplitude modulations. The TPAF has been shown in computer simulations to provide improved clutter discrimination over that of the BPOF. 4.5 In computer simulations

reduction in distortion sensitivity over that of the simple BPOF, while maintaining acceptable discrimination capability. The potential for implementing the TPAF in a real-time optical processor was demonstrated recently. Using the magneto-optic SLM (MOSLM) developed by Litton Data Systems and manufactured by Semetex Corporation. In prior work, the zero implitude modulation state of the TPAF was used to create a deblock in a BPOF with resulting improvement of 1.55 dB in measured experimental correlation signal-to-noise over the simple BPOF. The implementation of more complex ternary patterns has been shown in simulations to generate even greater improvements in correlation peak-to-clutter (P/C) ratios.

2. BACKGROUND

Discrete-valued spatial filter formulations have been developed to exploit the limited modulation capabilities of most commercially available SLMs. 10 The BPOF is a discrete-valued filter that encodes two phase modulation values, 0 and π in phase or 1 and -1 in amplitude. In an experimental optical correlator using a MOSLM at the filter plane, Psaltis et. al 11 demonstrated that thresholding on the real components of the Fourier transform produced a BPOF capable of pattern recognition. Horner and Leger 12 showed with computer simulations that BPOFs formulated by thresholding on the imaginary component of the Fourier transform also were capable of pattern recognition. The choices of thresholding on the imaginary or real component of the reference image's Fourier transform represent the two exremes of a continuum of thresholding-line angle (TLA) values. The TLA is defined as the angle between the thresholding line and the imaginary axis of the complex plane and it determines the relative amounts of even and odd reference image components encoded in the BPOF. A BPOF created with TLA = 0 is derived from the even component of the reference image. The BPOF formulation can be defined as follows:

$$F'(u,v) = \begin{cases} +1 & \text{Re}[F(u,v)]\cos\theta - \text{Im}[F(u,v)]\sin\theta \ge 0, \\ -1 & \text{otherwise}, \end{cases}$$
 (1)

where θ is the TLA. F'(u,v) is the BPOF pattern, and F(u,v) is the Fourier transform of the reference pattern.

Favorable characteristics of the BPOF include the ease of implementation in real-time SLMs, the high optical efficiency of a phase-only filter, the minimization of memory required to store the filter pattern, and a sharper correlation peak than a complex matched filter. Possible disadvantages of the BPOF relative to the matched filter include increased distortion sensitivity, reduced signal-to-noise ratio in the correlation response for classical input-noise cases, and the lack of amplitude modulation capability. The design and optimization of BPOFs have received considerable attention. ^{10,14-17}

The ternary phase-amplitude filter was developed to improve discrimination performance over that of the BPOF and related smart filters by providing the capability to prevent the energy of certain input spatial frequencies from reaching the output plane of the correlator. Since BPOFs modulate only the phase of the input spectrum, the energy of all the input spatial frequency components, including nontargets, is passed to the output plane, where it inevitably limits correlation signal-to-noise performance, i.e., discrimination against nontargets. The TPAF couples the BPOF's target recognition capability with the ability




Fig. 1. Experimental TPAF correlator system.

to selectively eliminate input spatial frequencies most closely related to nontargets, thus providing improved nontarget discrimination capability. The TPAF formulation can be described as the product of a BPOF with a binary amplitude mask (BAM). In general, BAMs are nontrivial patterns formulated with the aid of a procedure such as the transform ratio (TR) technique $\frac{1}{2}$. The TR procedure uses the ratio R(u,v) of the representative power spectral density pattern of target $|f(u,v)|^2$ and nontarget $|O(u,v)|^2$ patterns to determine the zero amplitude modulation states of the TPAF as follows:

$$F(u,v) = \begin{cases} \pm 1 \text{ (per BPOF)} & R(u,v) \ge T \\ 0 & \text{otherwise} \end{cases}$$
 (2)

In applying this technique it is important to vary the threshold (T) to optimize filter performance for the application being addressed.

3. OPTICAL CORRELATOR SYSTEM

The experimental optical correlator uses 48×48 MOSLMs at both input and Fourier planes and is depicted in Fig. 1. Polarizer P_1 following the input SLM is oriented in the binary amplitude modulation mode, while polarizer P_2 following the filter SLM is oriented in the binary phase modulation mode. The design details of this system have been reported and only will be summarized here. The classical or "4f" optical train involves two cascaded lenses positioned to perform Fourier transform operations on the input image and on the product of the input's Fourier transform and the spatial filter, respectively. To reduce the total correlator length, we use positive-negative lens pairs to perform the Fourier transform operations. The positioning of the input plane (SLM₁) at the first lens and the use of the positive-negative lens pair LP₁ to perform the Fourier transform operation results in a properly scaled Fourier transform with a quadratic

phase error that can be compensated by a positive lens (L_2) positioned near (ideally at) the Fourier plane occupied by SLM_2 . ^{2.18}

The optical train uses identical, mirror-image positive-negative lens pairs composed of lenses with focal lengths of 100 and - 50 mm. The phase correction lens L₂ placed adjacent to the filter SLM has a focal length of 292 mm. The length of the entire optical train for the 48×48 MOSLM (0.127 mm pixel spacing) system, excluding the HeNe laser source, is 1.22 m. This represents a factor of four reduction in length over the classical 4f configuration. The illumination subsystem, not shown in Fig. 1. includes a HeNe laser producing a 24 mW beam, which is spatially filtered with a 5× objective and 25 µm pinhole and collimated with a 160 mm achromat lens. The data acquisition subsystem includes an RCA TC-1030 silicon intensifier target video camera and an Imaging Technology Inc. Series 100 real-time digital image processing board in an NEC APC II (IBM PC AT compatible) computer. The signal is digitized at 8 bits per pixel and displayed with a video resolution capability of 640 horizontal and 480 vertical pixels. Detailed measurements indicate that the gamma of the video camera is approximately 0.7, and this is compensated for with a postprocessing operation.

The implementation of TPAFs in a real-time correlator reguires an SLM that is capable of simultaneous and independent binary phase and binary amplitude modulations. This modulation characteristic has been demonstrated with the MOSLM.7 The zero-state binary amplitude modulation capability of the MOSLM is realized by programming individual pixels into the neutral or mixed magnetization state in which about half of the pixel's active area is flipped from its original magnetic state to the opposite polarity. With the following polarizer set to the binary phase angle, these subpixel areas compose a random binary phase diffraction grating with characteristic spatial frequency. by definition, greater than the information bandwidth of the correlator (which is limited to the Nyquist bandwidth of the sampling associated with pixelation). Most of the light transmitted through the neutral state pixels is believed to be diffracted at angles larger than that associated with the coherent bandpass of the system as defined by the MOSLM's on-center pixel spacing, thus falling outside the useful correlation plane boundaries. Measurements on both 48×48 and 128×128 pixel devices. using coherent imaging with Nyquist bandwidth limits, have shown that the transmission of the mixed-magnetization zero state is approximately 3% in intensity relative to the phase-state transmission. This transmission, which we will call "leakage," is probably caused by the areas of the two magnetic domains in a pixel being unequal. This leakage is proportional to the difference in the two areas and involves a significant component at lower spatial frequencies falling within the Nyquist bandpass of the optical correlator, i.e., it will fall within the output plane used for correlation. The leakage component has the amplitude polarity (phase) corresponding to the domain having larger area in the pixel.

4. RESULTS

Computer-generated binary inputs and references, see Fig. 2, were used to demonstrate the implementation of nontrivial TR-TPAFs in a 48×48 MOSLM correlator. These patterns were combined in pairs to create input scenes for correlation. The X-O and E-F inputs represent character recognition, while the F1-F2 input represents target recognition of two different air-




Fig. 2. Computer-generated binary images used in studies.

planes. The first pattern of the input name (e.g., X in the case of the X-O input) is located in the upper left quadrant of the input scene, while the second pattern of the input is located in the lower right quadrant. Filters were based on centered reference patterns and a TLA of 0° (cosine BPOF).

4.1. Results compared to simulations using ideal zero state in the filter

Initial TPAF computer simulations were performed using a model that assumed ideal (opaque) zero-state amplitude modulation. The TPAF and BPOF results of the experimental and simulated studies are shown in Table 1.

For all five cases (corresponding to distinct reference patterns) there is generally good agreement between the simulated and experimental results. This is important because it indicates that analyzing correlation performance by computer simulation is reasonable, thus validating a convenient design tool. In addition. the table values show that for all five cases the implementation of the TPAF provided significantly improved discrimination over that of the BPOF. Notable improvements include those using the F2 and the E references correlated with the F1-F2 and the E-F inputs, respectively. The Horner efficiency values $\eta_{H,E}$ are based on the following definition: the quotient of the peak correlation pixel energy and the total input energy of the reference (i.e., target) pattern in the input pattern. Use of only the input target's energy as the base for efficiency is justified on the basis that it is inappropriate to penalize the filter's efficiency because clutter is present in the input, as would be done using the original definition. 19 Correlation P/C is defined as the ratio of the highest intensity value in the correlation peak resulting from the target to the highest intensity attributed to a nontarget pattern in the input and is expressed in decibel units.

An example of a significant improvement in discrimination is that of using the F2 target as the reference when correlating with the F1-F2 input. A TR-TPAF in which 58.4% of the filter elements were set to the zero state provided an experimental P/C improvement of nearly 3 dB, from 2.34 dB for a cosine BPOF to 5.13 dB. A typical technique used for discriminating between targets and nontargets is passing a constant threshold over the detected correlation plane at 3 dB down from the maximum peak value; using that technique the BPOF recognized both input targets to be the reference F2, while the TR-TPAF was able to discriminate between the two targets by 2.13 dB




Fig. 3. Relative intensity plots of representative correlation results from simulation (left) and experiment (right) for the E-F input and E reference target. (a) cosine BPOF (TLA=0), (b) TR-TPAF with 67.0% of the elements blocked.

TABLE I. TPAF and BPOF: results.

Reference	% Zeroed	Experimental	ideal Model	nn s (ideal)	Ternary Leak Model	
		P C (dB)	P.C (48)	*	P'C (dB)	
-	1				1	
	BPOF	7 66	7.97	30.9	ĺ	
	56.4	9 26	12.4	10.8	9.63	
0		j	İ	ł	t	
	BPOF	7 38	7.77	29.6	}	
	63.5	8.59	12.2	4.30	10.5	
£3			1			
	BPOF	2.34	2.58	12.0		
	35.7	2.66	3.06	9.20	2.90	
	46.0	4 03	4.42	5.34	4.33	
	50 7	4 28	4.45	4.61	4.25	
	58 4	5 13	6 86	2.64	6.17	
	70 1		6.33	1.30	8.90	
	75.7	-	8 60	0.85	8.62	
	52 4		8.95	0.49	7 75	
	99 1	· —	10.3	0.18	6.17	
	916	-	9.51	0.09	5.51	
F1			j			
	BPOF	5.27	5.34	10.8	į	
	. 41.6	5.44	6.00	6.25	6.68	
	54 8	6.20	6.73	3.66	8.60	
	58.4	7 86	8.78	2.38	9.86	
	64.3	9.39	9.90	1.35	10.4	
	72.7	_	11.0	0.50	9.83	
	80.7	ļ 	10.1	0.25	9.06	
	96.8		7.54	0.09	5.76	
	99.7		6.50	0.06	5 86	
E		1	İ	1	i	
	BPOF	1 60	1 73	15.2	ì	
	21.6	2.22	2.53	12.4	2.36	
	31 7	2 95	2.94	9.19	2 72	
	40.5	3.37	3 41	7.31	3 04	
	58 3	4 40	5.09	2.72	4 49	
	67.0	4 99	6.38	1 44	5.33	
	63.2	-	9 06	0 31	6.87	
	917		5 87	0.06	4 60	

A similar study in which the character E was used as the reference when correlating with the E-F input demonstrated a 3.39 dB improvement in P/C, from 1.60 dB for the cosine BPOF to 4.99 dB when a TR-TPAF with a blocking percentage of 67.0% was used. Again the TR-TPAF was able to discriminate between the E and the F, while the cosine BPOF could not. The amount of discrimination improvement obtained was similar for the other three references studied when the appropriate TR-TPAF was applied. Figure 3 qualitatively illustrates the improvements that can be achieved with TR-TPAFs in the form of linear-in-intentity plots. Note that the simulated TR-TPAF result plotted in this figure was generated from a model that assumes an opaque neutral pixel, while the experiment involves about 3% zero-state pixel leakage.

A study was performed to investigate the relationship between TR-TPAF blocking percentage (percent of filter elements that are opaque) and P/C and the general agreement between simulation and experiment over a wide range of blocking percentages. The blocking percentages reported in Table I were chosen in an attempt to sample a large portion of the blocking percentage range, encompassing a noticeable variation in P/C. The results of this study are summarized in Figs. 4 through 6, which also include computer model results for nonideal (i.e., "leaky") zero-state filter pixels to be discussed in a later section. The study was not performed for the X-O pattern since the discrimination performance of the simple cosine BPOF already was so good that adding the ability to make filter elements opaque had much less effect on the overall P/C performance. Figures 4




Fig. 4. P/C results from experiment and simulations with and without modeling of zero-state filter SLM leakage. The input was the F1-F2 pattern, while the reference pattern was F2.

through 6 indicate that as the TR-TPAF's blocking percentage increases, the P/C performance also increases, at least up to about 70% blocking. The figures also show good agreement between the computer simulation and experimental correlator performance. There is a practical limit to the increase in blocking percentage, however, because the Horner efficiency decreases as the blocking percentage increases. It was found that a blocking percentage of approximately 70% was the realistic maximum for generating correlation responses in our system. The decrease in Horner efficiency causes the system to become more susceptible to noise in the optical and detector systems, e.g., scattered light and co-erent noise, as well as other optical system imperfections such as imperfect alignment or nonideal SLM performance. This is seen in the results at higher blocking percentages where there is less agreement between the ideal (opaque zerostate) simulation and experiment. The most likely cause for this increase in hisagreement is that the simulation model does not take into account the filter SLM zero-state leakage discussed earlier. This effect should become much more prominent at high blocking percentages since a larger percentage of the total energy reaching the detected correlation plane is not predicted by the ideal model

4.2. Mode ing zero-state leakage

Additional simulation studies were performed with MOSLM zero-state it akage included in the computer model. An amplitude transmittance of 0.17 (corresponding to 3% intensity leakage, as measured) was used. The results of including leakage can be seen in Table I; the model including the ternary leakage phenomenon provides improved agreement with the experimental P/C for four out of the five cases studied. The improved simulation model is still imperfect; it does not include effects such as interference due to uncoated optical surfaces, detailed SLM pixel structure, etc. These imperfections, which undoubtedly degrade correlation performance, are less significant at low blocking percentages but tend to become more important as the Homer efficiencies decrease as they do for higher blocking percentages. This may belp explain why the agreement of the P/C between significal and experiment with F1 as the reference was not improved by including the leakage model; for this case Horner efficiencies average much lower for the same percentage




Fig. 5. Same as Fig. 4 except the input was the F1-F2 pettern, while the reference pattern was F1.




Fig. 6. Same as Fig. 4 except the input was the E-F pettern, while the reference pettern was E.

blocked than for the other references studied. The accuracy of the simulation model including leakage in predicting experimental P/C can be judged from results also plotted in Figs. 4 through 6.

4.3. Effects of inverting filter phase states relative to leakage

The polarity of filter zero-state leakage relative to the binary phase pattern may be switched by inverting the BPOF pattern programmed into the modulator. For the cases reported in previous sections, the leakage has a polarity identical to that of the fully saturated pixels modulating the +1 value of the TPAF. In an optical correlator using a TPAF, this leakage generates a spurious signal superimposed on the desired correlation amplitude pattern. The magnitude of this component will increase with increasing numbers of pixels blocked (set to the zero-amplitude state) in the TPAF and with higher leakage values. If a large proportion of filter pixels are set to the blocking state, the spurious output pattern component due to leakage is expected to resemble an inverted image of the input pattern. The "leaked" component interferes coherently with the correlation term and

TABLE II. TPAF phase reversal effects: results.

Ref	becreS F	Original	Original	Reversed	Reversed	Original	Reverse
		Experimental	Simulated	Experimental	Simulated	74 6	7# E
		P (* 48)	P (* 181	P C (4B)	P.C. 48)		
FI							
	BPOF	-	5.34		5 34	10 8	9 41
	41.6	5 44	6 00	6 50	6 96	5 24	4.73
	34.9	6 30	6 °3	. 28	9.78	2 63	2 14
	16.4	. 4	4 "4	: 12	10 5	1 73	1 25
	54 3	4 39	9 90		5 07	1 14	0 64
	72 7		110		4.14	0 50	0 16
	•0 7	-	10 1	-	3 33	0 25	0 13
Ε.							
	31.7		2 94		4 94	9 18	7 09
	40 3	••	3 41	-	5 44	7 31	5 20
	58.3	~-	5 09		7 29	2.72	1.74
	67.0	~	6 38	- .	3,94	1 44	0 61
	43 2	~	9 06		1 16	0 31	0 07
	91 7		3 47		1 31	0.06	0.05

produces effects that will probably be deleterious in most cases but will vary with the relative polarity (phase) of the leakage and binary phase modulation in the filter SLM.

The effect of reversing the phase of the fully saturated filter elements was investigated in computer simulations for the cases of the F1 and the E references when correlating with the F1-F2 and E-F inputs, respectively. The effect also was studied in the experimental optical correlator for the case of the F1. Simulations indicate that the phase reversal had the effect of improving the discrimination (P/C) performance at lower blocking percentages for both the F1 and E cases. However, when the blocking percentages exceed approximately 60% the effect of reversing the phase is to reduce the discrimination performance as seen in Table II. Reversing the phase also caused the Horner efficiency for all blocking percentages to decrease. Note that at the same blocking percentage (60%), reversing the phase increased the rate at which the Horner efficiency decreased with blocking percentage. The relationship between phase reversal, P/C. and Homer efficiency is not obvious, but it is clear that implementing higher blocking percentage TR-TPAFs will be more difficult and less productive with reversed polarity due to decreases in Horner efficiency and P/C, respectively.

TR-TPAFs with reversed phase values for the fully saturated filter elements were implemented in the experimental correlator for the F1 case. As seen in Table II and in Fig. 7, the agreement between simulation and experiment is good for the implemented original and phase reversed filters. The experimental responses produced with the phase reversed TR-TPAFs show an improvement in P/C over the original TR-TPAFs as predicted. The improvement in P/C is less obvious at the higher blocking percentages. As an example, the TR-TPAF with a blocking percentage of 64.3% generated a correlation response with nearly a 0 dB P/C value. A probable cause for this result is the dramatic decrease in Horner efficiency and the particular noise characteristics of our correlator system. However, in general, when optimizing the performance of a correlation system with MOSLMs that exhibit imperfect zero-amplitude modulation, the phase of the fully saturated filter elements relative to that of the leakage must be considered as a design parameter.

4.4. TLA effects

A final design parameter studied was the TLA used in BPOF and TR-TPAF formulation. Three different TLA values were investigated in this study: 0°, 90°, and 45°; the first value was used for all the results reported in previous sections. In this study the F1 target is used as the reference when correlating with the F1-F2 input. The blocking percentage used for computing the binary amplitude mask of the TR-TPAF was 41.6%. As indicated in Table III, the TR-TPAF using a TLA of 90° provided the best simulated P/C performance but also the lowest Horner efficiency. The agreement between simulation and experiment is good for the filters formulated from the 0° and 45° TLAs but is not nearly as good for the TLA of 90°. A qualitative comparison of the agreement between simulation and experiment is seen in Fig. 8. Much of the experimental noise in the correlation response is relatively unaffected by the change in TLA and appears more prominent at a TLA of 90° because the total energy in the correlation peak is much lower for that case. These simulations and the corresponding experimentation do not indicate that the Hartley filter (TLA = 45°) will always provide the best discrimination performance. It may be that the Hartley filter, in general. is a good choice because it encodes both even and odd reference. pattern information, but the creation of an optimum filter for a specific BPOF or TR-TPAF application must involve the consideration of the TLA as a design parameter, as has been reported. 13.15

5. SUMMARY AND CONCLUSIONS

In this work, BPOF and TR-TPAF formulations were investigated in both computer simulation and an optical correlator for several different in-class and out-of-class objects representing character and target recognition problems. Improvements in P/C of as much as 3.6 dB over corresponding BPOFs were obtained in the optical correlator with TR-TPAFs having blocking percentages as high as 67.0%. The degree of agreement between simulation and experiment was generally good but varied from case to case. Trends indicate that at low blocking percentages the agreement was very good, while agreement weakened at high blocking percentages. When the zero-state SLM pixel leakage of 3% in intensity was included in the simulation model, improved general agreement with experimental results was exhibited.

Studies were performed to determine the effect of reversing the phase of the saturated filter elements relative to the leakage through the zero-state elements. The investigations indicated that at low blocking percentages, reversing the phase provided improved discrimination performance and reduced Homer efficiency. At high blocking percentages the degree of reduction in Homer efficiency due to reversing the phase increases, which coincides with a trend of reduced discrimination performance in simulations. Discrimination is further reduced in actual optical systems, where the dynamic range of the input illumination is always limited.

Finally, the effect of the TLA used in setting filter phase modulation on correlator performance was investigated. For the cases studied, TLA angles of 45° and 90° provided better discrimination performance than a TLA of 0°, at the expense of Horner efficiency. In an optical correlator this relationship may have practical significance in terms of optimizing the practically achievable discrimination performance improvement with the constraints of limited source and detector dynamic ranges.




Fig. 7. Relative intensity plots of correlation results from simulations (left) and experiments (right) using phase reversed TPAFs. The input scene was F1-F2 and the reference pattern was F1. Three values of blocking percentage are shown: (a) 41.6%, (b) 54.8%, and (c) 58.4%.

These results have demonstrated the experimental achievement of significant improvements in correlation performance, in the form of increased discrimination against nontarget input patterns, which is easily realized with ternary phase-amplitude filter formulations implemented with commercially available MOSLMs. The use of computer simulations to predict and analyze correlation performance and assess filter designs was validated by the good agreement exhibited between such simulations and matching experimental results. Future work will address the experimental implementation of distortion-invariant TPAF formulations already showing promise in simulations

6. ACKNOWLEDGMENTS

This work was performed under Air Force contract F19628-87-C-0073 and sponsored by the Rome Air Development Center. RADC ESOP. Hanscom AFB. Mass.

TABLE M. TLA effects: results.

Reference	₹ Zero	TLA	Experimental	Simulated	~ * £
		ideg	P C (dB)	P ∈ dB	
Fl	41.6				
		0 00	6 39	6 00	5-24
		45.0	6 30	7.08	4.60
		90 0	- 05	9 29	3 🛺

7. REFERENCES

- A. VanderLugt, "Signal detection by complex spatial filtering. IEEE Trans. Inform. Theory 10(4), 139–145 (1964).
 D. Flannery, J. Loomis, M. Milkovich, and P. Keller. Application of binary phase-only correlation to machine vision," Opt. Eng. 27(4), 3484–320.

EXPERIMENTAL INVESTIGATION OF TRANSFORM RATIO TERNARY PHASE AMPLITUDE FUTERS FOR IMPROVED CORRELATION DISCRIMINATION




Fig. 8. Relative intensity plots of correlation results from simulation (left) and experiment (right) using F1 as the reference pettern and the F1-F2 input scene, for three values of TLA: (a) 0°. (b) 45°, and (c) 90°.

- S. Mills and W. Ross, "Dynamic magneto-optic correlator: real-time operation." in Acousto-Optic, Electro-Optic, and Magneto-Optic Devices and Applications. J. Lucero. ed., Proc. SPIE 753, 54-63 (1987).
 D. Flannery, J. Loomis, and M. Milkovich, "Transform-ratio ternary phase-
- amplitude filter formulation for improved correlation discrimination, Appl. Opt. 27(10), 4079–4083 (1988).
- M. Milkovich. D. Flannery, and J. Loomis, "Transform-ratio termary phase-amplitude filter formulations for character recognition," Opt. Eng. 28(5), 487–493 (1989)
- 6 D. Flannery, J. Loomus, and M. Milkovich, "New formulations for discrete
- D Flannery, J. Loomis, and M. Milkovich, "New formulations for discreasivalued correlation filters," in *Digital and Optical Shape Representation and Pattern Recognition*. R. Juday, ed., Proc. SPIE 938, 206–211 (1988). B. Kast, M. Giles, S. Lindell, and D. Flannery, "Implementation of ternary phase-amplitude filters using a magneso-optic spatial light modulator," Appl. Opt. 28(3), 1044–1046 (1989).

 S. Lindell and D. Flannery, "Ternary phase-amplitude filters for character recognition," in *Optical Information Processing Systems and Architectures*,
- B. Javidi, ed., Proc. SPIE 1151, 174-182 (1990).
 W. Ross, D. Psaltis, and R. Anderson, "Two-dimensional magneto-optic spatial light modulator for signal processing," in Real Time Signal Processing V. Proc. SPIE 341, 191-198 (1982).

- D. L. Flannery and J. L. Horner, "Fourier optical signal processors." Proc IEEE 77, 1511-1527 (1989).
- D. Psaltis, E. Pack, S. Venkatesh, "Optical image correlation with a binary spatial light modulator," Opt. Eng. 23(6) 698-704 (1984).
 J. L. Horner and J. R. Leger, "Pattern recognition with binary phase-only control of the control of t
- L. Horner and J. K. Leger. "Pattern recognition with binary phase-only filters," Appl. Opt. 24(3), 609-611 (1985).
 D. Flannery, J. Loomus, and M. Milkovich, "Design elements of binary phase-only correlation filters," Appl. Opt. 27(10), 4231-4235 (1988).
 F. Dickey, K. Stalker and J. Mason, "Bandwidth considerations for binary phase-only felters," Appl. Opt. 37(3), 3811-3818 (1988).
- phase-only filters," Appl. Opt. 27(9), 3811–3818 (1988).

 15. B. V. K. Vijaya Kumar and Z. Bahri, "Efficient algorithm for designing a ternary valued filter yielding maximum signal to noise ratio." Appl. Opt. 28(5), 1919–1925 (1989).
- 16. M. Farn and J. Goodman, "Optimal binary phase-only matched filters." Appl. Opt. 27(11), 4431–4437 (1988).

 17. D. Cottrell, R. Lilly, J. Davis, and T. Day, "Optical correlator performance
- of binary phase-only filters using Fourier and Hartley transforms. Appl Opt. 26(10), 3755-3761 (1987).
- 18. J. Goodman, Introduction to Fourier Optics, McGraw, New York (1968)
 19. J. L. Horner, "Light utilization in optical correlators," Appl. Opt. 21(12). 4511-4514 (1982).

Proc. SPIE, Vol. 1297 Orlando, Florida 20 Apr 1990

APPENDIX I

Application of Ternary Phase-amplitude Correlation Filters to LADAR Range Images

David L. Flannery and William B. Hahn The University of Dayton Research Institute Dayton, Ohio 45469

and

Edward R. Washwell Lockheed Missiles and Space Company Astronautics Division Sunnyvale, California 90488

ABSTRACT

Ternary Phase-amplitude Filters (TPAFs) used in a real-time hybrid (optical/electronic) correlator system comprise a promising pattern recognition approach with potential for near-term practical applications. Range images obtained from LADAR sensors present unique problems due to their particular signal and noise characteristics. We report computer simulation results of the application of both known and new TPAF formulations to the problem of target recognition on actual LADAR images. Binary input images suitable for input using magneto-optic spatial light modulators were generated by a simple preprocessing step which seems particularly suitable for these images. Experimental results verifying the simulated smart filter performance are presented.

1. Introduction and Background

Optical correlation using discrete-level filters has potential for near-term pattern recognition system implementations due to a combination of recent device technology and filter formulation advances [1-7]. Magneto-optic spatial light modulators (MOSLMs) have been shown capable of implementing both binary phase-only filters (BPOFs, with modulation levels of -1 and 1) [8,9] and ternary phase-amplitude filters (TPAFs, with modulation levels of -1, 0, and 1) [10,11], and both discrete-level filter types have demonstrated good correlation performance in theory and practice for a variety of pattern recognition problems [5,7,9,11,12,13]. The commercial availability of the relatively well developed MOSLM devices capable of frame rates of several hundred per second (from Semetex Corp., Torrance, California) supports the viability of practical near term system applications.

The well known distortion sensitivity of correlation must be confronted to solve any pattern recognition problem of practical complexity, even if hundreds of distinct reference images (filters) can be sequenced through the correlator each second, as with MOSLM-implemented BPOFs or TPAFs. This inspires efforts to formulate "smart" filters with widened in-class (e.g., target) response but retaining adequate out-of-class discrimination (e.g., nontarget pattern rejection) compared with a fiver based on a single view of the target. The formulation of smart filters is challenging because (1) typical problems are complex multi-class discrimination problems with ambiguous, or at best very complex, application-specific evaluation metrics (involving not only correlation peak strengths but shape factors), (2) finding the best filter comprises a multivariate optimization problem involving dimensionality equal to the number of resolvable filter elements (e.g., 16,384), and (3) for discrete level filters, constrained modulation levels must be implicit to the optimisation (i.e., a filter formulated for general complex modulation probably will not perform well when simply converted by thresholding to discrete modulation levels) [14].

A number of smart filter formulations for general complex filter modulation and, more recently, for discrete and constrained modulation levels have been reported, and a recent review may be found in reference [14]. In this paper we report a new TPAF formulation and compare its performance on LADAR (laser radar) range imagery to several other formulations. The LADAR images and preprocessing to obtain binary correlation input images will




Figure 1: Example of LADAR relative range scene processed for viewing

be discussed in the next section. Section 3 will present the distortion-invariant problem description chosen to close effort. In ensuing sections, the pertinent filter types, including the new filter formulation, will be discussed and computer simulation results exploring the performance of these filters on the test problem will be presented. Introduce perimental results verifying successful implementation of the new smart filter are then presented.

2. LADAR Range Images and Preprocessing

Range images obtained by processing of LADAR sensor data have a large dynamic range and particular noise characteristics which pose problems regarding their direct introduction into an optical correlator as analog amplitude modulation. This section discusses these characteristics and the level-slice binarization technique applied to produce binary image inputs suitable for the correlation process.

Laser range images processed in this effort were furnished by the Army Night Vision Laboratory and were relected from the March 1989 Field Test results [15]. The "relative range" images were used, with each pixel value a 16-bit integer representing relative range, thus defining a potential dynamic range in excess of 10,000-to-1. Some preprocessing is desirable simply for visual presentation of such images. The method we use involves handing the pixel value by 10 and taking the residue modulo 256. This results in an image such as the one shown in Figure 1, which contains range ambiguities introduced by the preprocessing algorithm, but also shows targets to good advantage. Targets typically span a range interval only a fraction of that corresponding to the modulo-256 induced ambiguity and thus usually are rendered in a relatively constant shade of gray, as seen in the figure. Histograms of pixel values comprising the background, usually terrain and trees, show a very uniform distribution over intensity. We have not performed statistical analysis of the background noise but it gives a strong visual impression of being largely spatially incoherent, except for some coherent features obviously related to terrain topography interacting with the range cycling induced by the modulo-256 preprocessing.




Figure 2: Binary level-slice point transfer function.

The modulo-256 preprocessing mentioned reduces the images to a dynamic range compatible with coherent optical correlation. However, correlation simulations performed with such images (e.g., Figure 1) demonstrated two problems: (1) the signal-to-(background) noise of the inputs is so poor that correlation even with simple BPOF references exhibited marginal peak-to-clutter, and (2) even this performance was obtained only when the target happened to fall at a high intensity level after preprocessing, while the opposite result (dark target) is equally probable. Thus a simple preprocessing step to eliminate or reduce these effects and also produce a binary image, facilitating image input via MOSLM devices was sought.

A binary level-slice technique, named "bin-slice," was developed. The point transfer characteristic defining this algorithm, as applied to the raw 16-bit data, is depicted in Figure 2, but it can be easily understood when viewed as being applied to modulo-256 preprocessed images. For example, a typical "bin-slice" image might be obtained by assigning maximum (255) pixel value to those input pixels having values ranging from 0 to 63, and zero to all others. Figure 3 shows two examples of applying this process to the same starting image. Note that the same target may appear bright or dark depending on the slice-levels chosen. Binary input images having a bright target, such as the truck in Figure 3 (a), provide good correlation performance using simple BPOF references, partially because of a reduction in background noise implicit to the level-slice process performed with quarter-maximum slice interval, and also because the target is at maximum amplitude. This sets the stage for design of distortion invariant filters, which invariably involve a peak-to-clutter tradeoff relative to simple BPOFs.

The strategy for using bin-sliced input images is as follows: the rapid sequential correlation capability inherent to a real-time optical correlator is used to advantage by processing a number of bin-sliced images derived from each distinct raw input image. Based on the results we report later in this paper, slice intervals of one quarter of maximum pixel value provide adequate noise reduction; thus correlation inputs derived from four contiguous slices would span the entire input dynamic range, ensuring a bright rendition of the target in one of the four input patterns. To handle the cases in which a target's pixel values are split across two slice intervals, six overlapping slices could be used. Thus six correlations (six different inputs, one filter) are needed to process each raw input image against a particular filter. This is not a daunting prospect given the capability to perform several hundred or more correlations per second as expected in the near future. We note that the operations required to perform bin-slice preprocessing are simple and should be amenable to real-time implementation with minimal hardware in either analog or digital form.

Quarter-amplitude bin-sliced images were used for all the studies reported here.




Figure 3: Examples of bin-slice operation applied to Figure 1, (a) bright truck case, (b) bright tank case.

3. Distortion-invariant Filter Design Problem

This section describes the smart filter design problem based on LADAR range images chosen for the reported effort, which comprised an initial attempt in terms of both the type of images processed and the new type of filter applied.

The 128-by-128 pixel LADAR images shown in Figures 1 and 3 contain broadside views of two potential targets: a truck at lower right and a tank at lower left. Other images contained broadside views of a helicopter, with a different scale. For these initial studies we chose the helicopter as the target and superimposed it above the truck and tank. It was isolated, scaled down to matching resolution, and superimposed by replacement on the scene of Figure 3 (a), yielding binary input scenes such as Figure 4, in which the helicopter is rotated five degrees clockwise relative to nominal attitude. The filter design problem was to formulate a filter which provided acceptable correlation performance over a ±10 degree rotation range of the helicopter target in scenes with the type of background shown.

Correlation performance must be defined. Peaks were defined using a 5-by-5 pixel box. The desired (target) correlation peak was required to be located within such a box centered on the ideal location to be counted a correct response. Figure 5 shows this box as well as others to be discussed superimposed on the scene of Figure 4. Peaks were recognized as distinct only if they fell outside the 5-by-5 box(es) of nearby higher peaks. Correlation efficiency was characterized by a modified Horner efficiency [16] defined as the energy in the target peak pixel divided by the total target energy in the input plane. A side-lobe box was defined in the target region as shown in Figure 5. (the larger box centered on the helicopter). Peaks detected in this region were declared to be target-dependent sidelobes rather than clutter. A false target box is also shown in the figure, centered on the truck image. Peaks located in that box were declared to be responses to the truck, and were used to form a metric for discrimination against that false target – simply the ratio (expressed in Decibels) of the properly located target peak energy to the highest peak energy located in the truck box. Peaks not classified into one of the boxes discussed were classified as clutter responses and were used to form the peak-to-clutter metric in a manner similar to that just discussed.



Figure 4: Example of binary input to correlation. Helicopter at five degrees clockwise rotation




Figure 5: Target, Sidelobe and Truck boxes used for classifying correlation peaks

This scene involves a dark tank which yielded no significant correlation responses; thus a tank discrimination box was not used.

4. Discussion of Filter Types

To provide a basis for comparative evaluation of the new filter formulation, a number of other formulations were investigated on the same correlation problem. These will be discussed in the following subsections and a description of the new formulation will be given.

4.1 Simple BPOF

The simple BPOF is a baseline for comparison of more advanced BPOF and/or TPAF formulations. Typically it provides more than adequate correlation performance for a perfectly matched input but exhibits unacceptable distortion sensitivity, although its performance even for a matched input can be seriously degraded by sufficient amounts of noise added to the input scene. The BPOF used in this study was based on a threshold line angle (TLA) of 45 degrees, which defines the Hartley BPOF [17,18]; this has proven to be a good choice for most patterns. This BFOF exhibited (in simulations) 6.2 dB P/C and a correlation efficiency of 2.8% for zero degrees of target rotation, and both metrics degraded by 3 dB for about ± 1.5 degrees target rotation. Clearly it falls far short of the design goal, but the rotation sensitivities were the basis for choosing training set intervals of two degrees for use in formulating the other filters.

4.2 Composite BPOF

A simple and intuitive way to formulate a distortion-invariant BPOF is to base it on a composite transform, i.e., the Fourier transform of a superposition of views representing a training set spanning the distortion being addressed. In this study, an equally-weighted composite of 11 rotations of the helicopter at two-degree intervals was used, thus spanning the desired ±10 degree interval. We note that iterative adjustment of the composite weights to achieve a desired response distribution over the training set (e.g. equal correlation peaks) amounts to the Jared-and-Ennis (SDF-BPOF formulation [19]. We did not apply this refinement to our composite BPOF but expect that it would indeed accomplish its goal if it were applied.

4.3 Composite Transform-ratio TPAF

A TPAF may be viewed as the product of a BPOF pattern and a binary amplitude mask (BAM). The transform-ratio (TR) technique provides a method of generating the BAM based on thresholding the ratio of power spectra of the in-class and out-of-class patterns [6]. An out-of-class composite power spectrum was formed by the sum of the spectra of the two background scenes shown in Figure 3. In general more robust combinations would be preferred to properly statistically represent out-of-class patterns. The in-class spectrum was that of the helicopter composite image mentioned in the previous subsection, and the BPOF pattern was the composite BPOF also described in the previous subsection. Several choices of TR threshold were tried and the results to be presented (the best obtained) used a threshold which corresponded to 83.03% of the filter pixels being zeroed.

4.4 Filter-TPAF Formulation

The filter-TPAF (fTPAF) is an extension of the Jared-and-Ennis fSDF-BPOF [19] to incorporate a BAM [13], and also may be viewed as the composite TR-TPAF described in the previous subsection with the superposition weights adjusted iteratively to achieve equal correlation peaks over the training set. Note that filter training is done with clean (sero background) images, so equal peak responses using actual inputs including backgrounds are not guaranteed. The BAM is determined independently of and prior to the filter design iteration, and in this case the BAM described in the previous subsection was used. The same 11 views of the helicopter mentioned in previous paragraphs were used as the training set.

4.5 A New Smart TPAF Formulation: the Metric-Sort

The new filter formulation, which we call "metric-sort" or MS for brevity, is an iterative numerical process which attempts to maximize the classical SNR of a filter constrained to TPAF modulation values. A numerical procedure for solving this problem for the case of additive white noise has been given recently by Kumar and Bahri (20). Our treatment differs from theirs in at least two important ways: (1) we treat any noise spectrum, removing the need to assume white noise, and (2) whereas their algorithm was exact at a exhaustive (i.e., guaranteed to find the global maximum SNR) ours involves approximations and is subject to trapping to local SNR maxima, depending on an assumed starting point. The above comments apply to the case of a single target pattern with additive stochastic noise. In the work to be reported here we made the further ad hoc extensions of using a composite target image (defined in pravious paragraphs) and applying the filter to scenes in which noise was not additive (because target superposition was done on a replacement basis). In addition, our composite noise sample is suspected of not being truly stochastic. Nevertheless we obtained better results for the chosen correlation problem than with any of the other filters discussed above.

The classical SNR metric addressed by our MS procedure is:

$$SNR = \frac{|\int S(f)H(f)df|^2}{\int N(f)|H(f)|^2|df},$$
 (1)

where f is the (spatial) frequency variable, and S, H, and N are the signal Fourier transform, filter function, and noise spectrum, respectively. A brief description of our formulation algorithm follows:

- 1. Assume a starting TPAF pattern, for example a BPOF formed with a particular TLA, and calculate and store the corresponding SNR (Eq. 1).
- 2. For each filter pixel, calculate the differential SNR resulting from switching the pixel to each of the remaining two allowable filter modulation states (-1, 0, 1). Store the maximum of these two values and the associated new modulation state.
 - 3. Sort the pixels into descending order of the SNR differentials obtained in step 2.
- 4. For the top M values on the sorted list, make the associated pixel modulation state changes and update the SNR value. If a negative SNR differential is encountered, stop and go to step 6.
- 5. Go to step 2, except if an operator-selected maximum number of pixels to be blocked has been exceeded go to next step.
 - 6. Write the new TPAF filter to a file and end.

Several comments apply to this procedure. First, it clearly has the general nature of a gradient-descent or hill-climbing algorithm. Thus we expect it to be subject to trapping at local maxima and sensitivity to starting point assumptions. Second, an ideal implementation would require M=1, i.e., resorting the differentials after each single pixel is switched to a new modulation state. This would be extremely compute-intensive, and we suspect not necessary, provided M is very small compared to the number of filter elements, N. The fractional change in SNR in response to changing a single element (out of, for example, 16,384 in a 128-by-128 element filter) is very small. Typically we have used 0.05N as the maximum value for M. The reason for the exception in step 5 is that empirical evidence indicates that filters which optimize SNR are frequently deficient in other metrics such as peak height (Horner efficiency) or peak width. Wider peaks (i.e., poorer location estimates) and lower efficiency typically are associated with increased fractions of pixels blocked. Thus we allow for the possibility of stopping the iterations before convergence to preserve better overall performance, based on the filter designer's judgment.

In the work reported here, five different initial TPAFs were tried as starting points, comprising BPOFs formulated with TLAs of 0, 22.5, 45, 67.5, and 90 degrees. Best overall performance was obtained by stopping iterations when about 54% of filter elements had been switched to the zero-state. The five resulting filters showed slight variations but approximately equal overall performance for all TLAs except 90 degrees, which was noticeably inferior. The filter resulting from a starting TLA of 45 degrees was chosen as representative of best overall performance and was used in results to be reported.




Figure 6: Peak-to-clutter performance of four TPAF smart filters.

5. Simulation Results Comparing Filter Types

The four smart filters defined in the previous section were tested over the ± 10 degree target rotation training range but using one degree testing intervals to check interpolation between training set cases.

5.1 Peak-to-clutter and Horner Efficiency

Figure 6 presents the P/C performance of the filters, and Figure 7 presents the corresponding correlation peak energies. The helicopter target image contains 525 pixels of value unity, or a total energy of 525 units which is the basis for Horner efficiency. Thus full scale in Figure 7 corresponds to a Horner efficiency of 1.43%. The overall superiority of the MS filter in terms of combined P/C and Horner efficiency is apparent from examination of these two figures. Note that the fTPAF exhibits less "droop" at the end points of the rotation interval, as might be expected. We expect that iterative adjustment of the composite weights used to train the MS filter could reduce this droop in the same manner, or the present filter could be viewed as one providing excellent uniformity over a reduced interval, e.g., ±8 degrees.

5.2 Discrimination and Sidelobes

Discrimination against the truck and sidelobe levels were evaluated over the same testing interval.




Figure 7: Target correlation peak intensity of four TPAF smart filters.

The minimum truck discrimination and maximum sidelobe levels found over the central ±8 degree interval are tabulated:

Filter Type	Minimum Truck Rejection (dB)	Maximum Sidelobe (dB)
MS	3.7	-3.5
fTPAF	1.6	-1.9
Comp. TR-TPAF	2.2	-1.9
Comp. BPOF	0.8	-0.2

The MS filter is again generally superior to the other types.

5 3 Interpretation of Simulation Results

In these simulations, clearly the MS filter was significantly superior to the other three types of smart TPAFs tested, based on overall performance with respect to four metrics. The generality of this result is limited by the use of only one test background scene, Figure 3 (a), and also by the other specific characteristics of the defined problem. The scope of these tests also was limited in that exhaustive design iterations to obtain the best possible exemplar of each filter type were not performed. However, prior experience with the fTPAF and TR-TPAF filters indicates that any change in the percentage elements seroed involves a tradeoff between P/C and Horner efficiency, so that the likelihood of finding a better overall combination than the filters tested seems low. The possible variations of

MS filter formulation procedure have only been sparsely explored in this effort. It seems possible that other choices of filter starting point might yield substantially improved results.

The MS filter was tested using the other background scene, Figure 3 (b) and exhibited about the same P. C and Horner efficiency, but substantially reduced rejection of the (tank) false target (e.g., 1-2 dB). We conjecture this is because the tank looks more like a composite helicopter and/or because the tank image contains more energy.

In one case, TLA = 0, the MS filter was iterated to convergence, resulting in 92.5% of the elements being set to the zero-state. This filter exhibited only slightly better P/C than the MS filter used in the above tests, although uniformity over the ± 10 degree testing interval was significantly improved. These improvements were obtained at the expense of a 4.3 dB reduction in Horner efficiency, to 0.29%, or a 9.8 dB reduction from the efficiency of the BPOF. The practical impact of this depends on the specific correlator hardware implementation. It probably still compares favorably to the efficiency of a classical matched filter.

6. Experimental Correlations

We have obtained initial experimental results with the MS filter for the defined test problem, for helicopter rotations of -10, -5, 0, 5, and 10 degrees, on the background scene of Figure 3 (a). The experimental correlator used 128-by-128 Semetex MOSLMs at both input and Fourier planes, and used an optical configuration similar to that previously reported [8]. Normalized intensity plots of the simulation and experimental results for +5 degrees rotation are shown in Figure 8.

The P/C and sidelobe levels exhibited by the experimental results are tabulated with the simulation predictions:

Rotation Angle (degrees)	P/C (dB) Measured	Simulated	Sidelobe (dB) Measured	Simulated
-10	5.3	3.9	-2.3	-2.5
-5	5.6	5.5	-3.8	-4.6
0	5.9	6.4	-4.1	-5.1
5	6.6	5.5	-3.2	-5.5
10	5.2	2.7	-0.6	-2.0
Avg.	5.7	4.8	-2.8	-3.9

The tabulated and plotted experimental results are based on captured video images and have been corrected for a gamma value of 0.7, believed to be appropriate based on careful measurements performed on a different sample of the same model camera. If the results are not gamma-corrected, the dB measurements in the above table are reduced to 70% of the tabulated values, i.e., an average P/C of 4.0 dB which is still in reasonable agreement with the simulation average. Sidelobe levels are already higher than predicted by simulations and would become higher without gamma correction.

Regardless of possible questions concerning gamma correction, these results clearly enable easy target peak detection (as can be judged from the figure) and they represent excellent agreement between experiment and simulation, considering that the correlator has just been assembled and probably can be adjusted for improved performance.




Figure 4: Intensity plots of correlation response using metric-sort TPAF smart filter. (a) Simulation. (b) Expensional

7. Conclusions

A new TPAF smart correlation filter formulation, the metric-sort (MS) filter, has been defined and assessed by computer simulations on a distortion-invariant test problem based on LADAR range imagery. The MS filter exhibited overall performance superior to three other smart TPAFs, based on four metrics addressing light efficiency and discrimination.

A simple and easily implemented preprocessing technique called the "bin-slice" algorithm was developed for application to LADAR range images to yield binary images suitable for input into optical correlators using binary input SLMs. A strategy for using this technique in conjunction with a high frame rate correlator system was suggested.

Experimental results verifying the successful implementation of the developed MS filter were presented

The MS filter formulation is in its initial stages of development and application. Its design variations have only been sparsely investigated. We are investigating techniques to incorporate both specific nontarget shapes and peak-width constraints in the MS design process.

8. Acknowledgments

The investigation of LADAR range images was sponsored by Lockheed Missiles and Space Company, Astronautics Division, Sunnyvale, California.

The reported filter developments were sponsored by the U. S. Air Force under contract F19628-87-C-0073 sponsored by the Rome Air Development Center, RADC/ESOP, Hanscom AFB, Massachusetts.

The LADAR range images were furnished by the U.S. Army Night Vision Labs, C2NVEO/AMSEL-RD-NV-V. Ft. Belvoir, Virginia. We thank Mr. An Nettleton of that organization for expediting their release.

The major portion of the equipment comprising the experimental correlator system was provided by Martin Marietta Strategic Systems, Denver, Colorado.

9. REFERENCES

- [1] W. Ross, D. Psaltis, and R. Anderson, "Two-dimensional magneto-optic spatial light modulator for signal processing," SPIE Vol. 341, "Real-time Signal Processing V," San Diego, CA, p. 191, August 1982.
- [2] J. Davis and J. Waas, "Current status of the magneto-optic spatial light modulator," Proc. SPIE, Vol. 1151, (in press), San Diego, CA, August 1989.
- [3] D. Psaltis, E. Paek, and S. Venkatesh, "Optical Image Correlation with a Binary Spatial Light Modulator," Optical Engineering, pp. 698-704, Nov./Dec. 1984.
- [4] J.L. Horner, J.R. Leger, "Pattern Recognition with Binary Phase-Only Filters," Applied Optics, pp. 609-611, Vol. 24, 1 March 1985.
- [5] J. L. Horner and H. O. Bartelt, "Two-bit correlation," Applied Optics, Vol. 24, pp. 2889-2893, 15 September 1985.
- [6] D. Flannery, J. Loomis, and M. Milkovich, "Transform-ratio ternary phase-amplitude filter formulation for improved correlation discrimination," Applied Optics, Vol. 27, pp. 4079-83, 1 October 1988.

- [7] M. Milkovich, D. Flannery, and J. Loomis, "A study of Transform-ratio ternary phase-amplitude filter formulations for character recognition," Optical Engineering, Vol. 28, pp. 487-493, May 1989.
- [8] D. Flannery, A. Biernacki, J. Loomis, and S. Cartwright, "Real-time Coherent Correlator Using Binary Magnetooptic Spatial Light Modulators at Input and Fourier Planes," Applied Optics, Vol. 25, pp. 466, 1 February 1986.
- [9] D. Flannery, J. Loomis, M. Milkovich, and P. Keller, "Application of binary phase-only correlation to machine vision," Optical Engineering, Vol 27, pp. 309-320, April 1988.
- [10] B. Kast, M. Giles, S. Lindell, and D. Flannery, "Implementation of ternary phase-amplitude filters using a magneto-optic spatial light modulator," Applied Optics, Vol. 28, pp. 1044-46, 15 March 1989.
- [11] S. Lindell and D. Flannery, "Ternary phase-amplitude filters for character recognition," Proc. SPIE. Vol. 1151, pp. 174-182, August 1989.
- [12] K.H. Fielding and J.L. Horner, "Clutter Effects on Optical Correlators," Proc. SPIE, Vol. 1151, pp. 130-137, San Diego, CA, August 1989.
- [13] D. Flannery, J. Loomis, and M. Milkovich, "New formulations for discrete-valued correlation filters," Proc. SPIE, Vol. 938, pp. 206-211, April 1988.
- [14] D. L. Flannery and J. L. Horner, "Fourier Optical Signal Processors," Proc. IEEE, Vol. 77, pp. 1511-1527, October 1989.
- [15] Private Communication, J. E. Nettleton, U.S. Army Center for Night Vision, C2NVEO/AMSEL-RD-NV-V.
- [16] J. L. Horner, "Light utilization in optical correlators," Applied Optics, Vol. 21, pp. 4511-14, 15 December 1982.
- [17] D. Cottrell, R. Lilly, J. Davis, and T. Day, "Optical Correlator Performance of Binary Phase-Only Filters Using Fourier and Hartley Transforms," Applied Optics, pp. 3755-3761, 15 October 1987.
- [18] D. Flannery, J. Loomis, and M. Milkovich, "Design elements of binary phase-only correlation filters," Applied Optics, Vol. 27, pp. 4231-4235, 15 October 1988.
- [19] D. Jared and D. Ennis, "Inclusion of filter modulation in synthetic-discriminant function filters," Applied Optics, Vol. 28, pp. 232-239, 15 January 1989.
- [20] B. Kumar and Z. Bahri, "Efficient algorithm for designing a ternary valued filter yielding maximum signal to noise ratio," Applied Optics, Vol. 28, pp. 1919-1925, 15 May 1989.

Presented at AAAIC '89 Dayton, OH October, 1989

ON-LINE NEURAL ADAPTATION OF OPTICAL CORRELATION FILTERS

David L. Flannery
Steven C. Gustafson
Darren M. Simon
The University of Dayton Research Institute
Dayton, Ohio 45469

abstract

Hybrid (optical/electronic) adaptive correlator (HAC) systems based on rapid sequential optical correlation are receiving renewed interest for automatic target recognition (ATR) primarily as a result of recently developed real-time devices and filter formulations which make compact, rugged, high performance systems possible in the near term. These systems will attack the class discrimination problem through the combined power of high correlation throughput and "smart" filter formulations which reduce the number of filters needed to perform a recognition task in the presence of clutter. Neural network approaches have potential application in several roles as part of a HAC system. Of particular interest is the concept of on-line adaptation of the filter pattern to maintain near-optimum correlation performance despite changes in input scene noise and/or clutter and target distortions. This paper reports initial results of a simulation study investigating this concept. The HAC concept and supporting technology developments are reviewed to provide necessary context.

1 Introduction

The hybrid adaptive optical correlator (HAC), diagrammed in Figure 1, is a leading candidate to achieve an impressive existence proof for useful optical processors in pattern recognition applications such as automatic target recognition. Research and device development over the past several years, primarily the combination of magneto-optic spatial light modulators [1,2] and the binary phase-only filter (BPOF) [3,4] and ternary phase-amplitude filter (TPAF) [5] concepts have led to the laboratory demonstration of real-time correlators which prove the feasibility of hybrid correlation concepts [6.7.8]. Systems which can be constructed in the near term based on these concepts will be capable of performing correlations of 256-by-256 element patterns at rates up to 1000 correlations per second. Each correlation can involve a different input and reference pattern, the latter being stored in digital memory. Such storage is practical and efficient due to the reduced information content of BPOF filters; 1000 such filters will fit in four MB (megabytes) of memory, a rather modest amount with today's digital technology. Rapid access of BPOF patterns is possible because they are stored in the binary form directly useable by the filter SLM; no D/A or other type of conversion is required. In the HAC concept, Figure 1, this real-time optical correlation approach is the basis for the "rapid sequential optical correlator" module, while the remainder of the system is envisioned as electronic, thus defining a hybrid system.




Figure 1: Hybrid Adaptive Correlator Concept Diagram

Neural network approaches have potential application in the HAC architecture in at least three general (and somewhat overlapping) functional areas:

- Adaptive processing to control HAC operation,
- Post processing to interpret correlation output patterns,
- Adaptive on-line formulation or modification of filter patterns.

In this paper we report initial results of a computer simulation study aimed at assessing the feasibility of one of many possible concepts in the third area listed – adaptive filter formulation. We note that work addressing the second area has already been reported by researchers at the Air Force Institute of Technology [9].

Background on discrete-valued correlation filters (i.e., BPOFs and TPAFs) will be presented, followed by a description of the relatively simple problem selected for our initial investigation and the neural network approach used. Preliminary results of this ongoing effort will be presented and interpreted.

2 Discrete-valued correlation filters

We are concerned with filter patterns to be introduced (modulated, addressed) into the well known "4-F" coherent optical correlator configuration based on Fourier optics [10]. In this processor, which uses light amplitude as the analog variable, convolution (or

correlation) functions are achieved by multiplication of light patterns representing the Fourier transform of an input pattern by a modulation pattern representing the Fourier transform of a reference pattern (i.e. filter pattern), followed by another Fourier (or inverse Fourier) transform, each transform being accomplished with a single lens. For practical (real-time) operation in a HAC, the filter pattern must be introduced rapidly using a spatial light modulator (SLM) device. Until recently this was believed impractical because of the lack of a SLM capable of modulating patterns with complex values, (having independently varying real and imaginary, or phase and magnitude, components).

In 1984, attractive correlation performance was demonstrated, in simulations, for filters retaining only the phase component of the reference function transform [11] and, shortly thereafter, correlation filters quantized to only two phase values, 0 and π radians or modulation levels of -1 and +1, were shown to perform well [3,4]. Since magneto-optic SLMs (MOSLM) capable of rapid binary-phase modulation were commercially available, successful demonstrations of real-time BPOF correlators were soon reported [6,7], and the potential for the HAC was established.

Ternary phase-amplitude filters use modulation levels of -1, 0, and 1, i.e., a combination of binary phase and binary amplitude modulation which may be viewed as the cascade (multiplication) of a binary amplitude mask (BAM) with an existing BPOF pattern. Formulation of the BAM pattern based on the ratio of target and nontarget spectral energies can improve discrimination or correlation peak-to-clutter ratios as has been demonstrated in simulations [5] and experiments [8]. The experimental demonstration used an easily accessed third state of the MOSLM devices as the zero modulation state. Thus TPAF filters are a practical reality for use in the HAC.

A primary issue of HAC operation is handling input target distortions while maintaining good target-vs-nontarget discrimination. Correlation is notoriously distortion sensitive so that a single filter has no hope of performing well for any problem of practical complexity, such as automatic target recognition. The HAC architecture has an inherent strength in attacking this problem: the large bank of (precomputed) correlation filters which can be rapidly accessed. Also, each filter can be designed to optimize the inherent tradeoff between distortion invariance and discrimination for the particular application domain. Such filters are called "smart filters."

Smart filters for BPOF and TPAF encoding have been reported [12,13,14] and their use in the HAC defines a powerful generic pattern recognition approach. Nevertheless, the ability to adaptively reformulate the filter pattern so as to maintain good filter performance in response to changing input parameters, such as target distortions and noise/clutter variations, is very desirable since any system relying solely on a bank of precomputed filters can be overwhelmed when the complexity of the application exceeds some limit determined by system filter storage capacity and/or the computations required to select the proper filters from storage.

In the work reported here, the use of a neural network to adaptively select the bandpass characteristic of a TPAF is investigated.




Figure 2: Input scenes with noise levels of (a) 10%, (b) 20%, (c) 40% and (d) 60%.

3 Problem Definition

An initial test problem was defined using binary input scenes as shown in Figure 2. The target is a hypothetical airplane silhouette and the nontarget is random background noise. Simulations were performed on a 64-by-64 pixel format. To generate the background noise, pixels were randomly turned on (set to binary one) with each pixel independently having a certain probability of being on. This probability was expressed as a percentage and served as an indication of noise level. The noise level also represents the fraction of background pixels turned on. Noise levels of 0, 10, 20, 40, and 60% were used in these studies.

A BPOF was computed for the target using a threshold line angle of 45 degrees [15] and was used as the phase component of all the TPAF patterns investigated. The issue under investigation in these initial studies is the performance of a neural network in




Figure 3: Bandpass BAM rings defined in Fourier (filter) plane

establishing the BAM pattern of a TPAF so as to adapt the filter for good peak-to-clutter performance in response to variations of the input noise level.

In general BAM formulation requires the optimal setting of N^2 or 4096 binary values, implying a neural network having that many binary outputs, which is currently beyond practical consideration, even in simulations. To reduce problem complexity to a manageable, but still meaningful, level the domain of possible BAM patterns was parametrized using bandpass rings as depicted in Figure 3. In this figure, the square borders define the Nyquist-limited bandwidth of the processor, assumed to be defined by a 64-by-64 element pixelated input SLM. The four BAM rings are defined by radii having factor-of-two ratios, i.e., 4, 8, 16, and 32 pixels. Each ring is either on or off (all encompassed pixels set either to one or zero) and a complete BAM specification simplifies to a four-bit binary number.

A significant body of "designer wisdom" is available regarding the effects of spatial frequency bandpass tailoring on correlation performance both from the literature [16] and from observations of many simulations performed in this and other studies. The primary mechanisms of interest here are the effect of bandpass on two key correlation performance metrics: peak-to-clutter and peak width. In our studies peak-to-clutter is defined as the ratio of the energy (intensity) of the properly located maximum correlation peak pixel to the next highest pixel energy value located outside the correlation peak region. This region was arbitrarily defined as a 5-by-5 pixel area centered on the peak. In general, there is a bandpass which optimizes peak-to-clutter for a given target and type of input noise. This optimization has been analyzed for phase-only filters and Gaussian zero-mean additive white noise [16] but analytical solutions valid for general noise types are not available.

Correlation peak width is important because it determines the difficulty of peak localization, which translates directly to target location estimate accuracy. In general, peak width is inversely related to filter bandwidth via the uncertainty principle of the Fourier transform, i.e., narrow peaks require high bandwidth.

Filter design must deal with the intrinsic tradeoff resulting from the peak-to-clutter and peak width dependencies on filter bandwidth, i.e., the BAM pattern yielding the best peak-to-clutter will frequently not be the one yielding the narrowest peak. We have handled this in our studies by a simple constraint designed to ensure a reasonably narrow peak – we exclude from consideration BAM patterns in which only the first (smallest radius) ring is turned on. This is a design choice based on examination of the correlation peak widths corresponding to different bandpass choices, and a different choice might be made depending on the relative importances of peak-to-clutter and localization performance.

To summarize, the goal of the neural network is to provide four binary outputs which set the on/off status of the four bandpass rings of a BAM (see Figure 3) so as to optimize peak-to-clutter performance of the resulting TPAF in response to varying amounts of input noise, subject to a maximum peak width constraint which excludes one of the 16 possible BAM choices. (One other bam choice, all zeros, is excluded because it completely blocks the signal, leaving 14 allowed BAM patterns.)

4 Neural Network Approach

We now discuss the configuration of a neural network to attack the problem defined in the previous section. The first major issue is to define the information the network should use to make its decisions, i.e., to define the network inputs. There are many possible choices, but we have selected one which is both attractively simple and, we believe, logical, i.e., probably includes the requisite information to support a good network choice. The power spectral energy of an input scene averaged over each of the four BAM rings defined in Figure 3 comprises a set of four continuous valued features which are easily derived in real-time by a slight modification of the optical correlator architecture and which represent a crude power spectrum analysis of the background noise (assuming that the target subtends a small fraction of the total field of view). The required correlator modification involves simply inserting a beam splitter after the first Fourier transform lens to direct the optically derived transform onto a suitable detector array.

We note that noise level variations for the noise model described can be characterized by only one parameter. Why not make that single parameter the input to the neural network? If handling noise generated by the particular noise model described was the only goal of our work, this would be a valid criticism. However, an implicit assumption of our problem setup is that other types of noise can be handled by the same approach and that power spectral density features are good candidates as general descriptors of a wide class of noise types.

Thus the desired neural network has four continuous inputs and four binary outputs. We selected a three-layer perceptron network with 13 hidden-layer units, trained by back-propagation [17].

Training and (separate) test sets were defined for the five noise levels mentioned, including multiple instances of input noise samples at each level above zero. Correlations were performed using each of the 14 allowed BAM (network output) choices for each of

Table 1: Selected training set cases


Noise level	Band 1	Band 2	Band 3	Band 4	BAM pattern	Peak/Clutter
(% area)	energy	energy	energy	energy	(binary #)	(dB)
10	1.416	0.168	0.093	0.083	1111	11.2
10	1.370	0.175	0.090	0.085	1110	11.4
20	1.296	0.229	0.139	0.143	1111	10.5
20	1.294	0.238	0.122	0.146	1110	11.2
40	0.859	0.209	0.194	0.210	1100	7.3
40	0.863	0.314	0.184	0.206	1110	6.8
60	0.463	0.198	0.190	0.208	1100	5.5
60	0.495	0.256	0.184	0.209	1110	4.74

33 sample input scenes. For training, two BAM patterns were chosen as desired network outputs for each input case. Selection of the two BAMs was based on peak-to-clutter and statistical representation considerations. In all, 35 training set cases were defined, encompassing five noise levels, 17 different input scenes, and 35 desired BAM output patterns. The remaining 16 input scenes (plus the zero-noise case) were withheld as a test set. A selection of training set cases is given in Table 1, in which the variation of the four input spectral features with noise level is represented.

A complication which must be addressed in formulating a viable training set is that typically 4 or 5 of the 14 allowed BAMs will provide peak-to-clutter within two dB of the maximum value for a particular input scene. The straightforward approach is to represent every BAM within this acceptable near-optimum range in a training set case, which would result in a training set several times the size we used. Instead we selected training BAMs with the goal of fairly representing all the different near-optimum BAM patterns in the training set.

5 Preliminary Results

The network was trained for 1,000 cycles and peak-to-clutter performance was evaluated. Performance on the training and test tests is summarized in Figure 4. To provide a fair basis for evaluation, the peak-to-clutter performance of two fixed "designer's choice" BAM patterns is included in the figure. These are the full-bandwidth filter and the quarter-bandwidth filter, corresponding to all or only the innermost two BAM rings being turned on, respectively. Maximum, minimum, and average peak-to-clutter values are indicated by cross-hatching variations on each bar. Results with zero input noise are not plotted since "peak-to-clutter" is at best a questionable concept in this case.






Figure 4: Neural-adapted filter performance for (a) training set inputs and (b) test set inputs. Noise levels (percentages) are listed along left side. "ANS" denotes Artificial Neural System. Min, Avg, and Max values are located at the right end of their respective cross-hatched areas.

In general an experienced designer might expect the full bandwidth filter to perform best with little or no input noise and the reduced bandwidth filter to perform better with higher noise levels, and these trends are seen in the figure, although all filters perform well with only 10% input noise. The performance of the adaptive filter must be judged as highly acceptable (i.e., near optimum). It provides peak-to-clutter very close to the best fixed-BAM filter and/or substantially better than the worst fixed-BAM filter in every case. Thus it is providing near-optimum performance when compared with the two designer's choice filters.

It could be argued that the fixed-BAM quarter-bandwidth filter provides better overall peak-to-clutter performance than the adaptive filter which then has no apparent advantage. However, the adaptive filter provides full bandwidth when the noise level is low enough to do so without sacrificing peak-to-clutter, and this provides narrower peaks which are preferable for target location estimation purposes.

Although not conveyed by the data presented, examination of the correlations performed with all possible BAM patterns indicates that the full-bandwidth and quarter-bandwidth BAMs are optimum or near optimum in terms of peak-to-clutter for the 10% and 60% noise levels respectively, and that a gradual transition of bandwidth from full to quarter as noise increases is a near optimum adaptive behavior. Examination of the BAM patterns output by the network indicates that it approximates this behavior.

Training was continued to 10,000 cycles with the same training set, resulting in few changes in the BAM pattern chosen for a given input, and no net improvement of overall filter performance.

A key issue is whether this adaptive filter approach can be successfully extended to different forms of noise. Figure 5 gives examples of input scenes containing different levels of background clutter generated by an Ad Hoc model which randomly positions noise clumps of randomly variable size. We first wondered how well the network trained only using the original type of noise would perform on the new type of noise. We found that the adaptive filter performed reasonably well, providing peak-to-clutter ranging from 8.7 to 2.8 dB over the noise range of 10% to 60% (area coverage) compared with the absolute-best BAMs (chosen differently for each noise level), whose performance ranged from 10.99 to 3.98 dB.

6 Conclusion

These initial simulation results indicate promise that neural networks can provide useful adaptation of TPAF correlation filters, at least for the relatively simple problem addressed. These preliminary results are based on only a sparse exploration of critical parameters such as number of hidden units, training set size and representation, and performance variation with number of training cycles. We plan to extend the investigation to encompass different noise/clutter models (e.g., Figure 5) and a more robust exploration of the important network parameters.

The problem of adaptation to follow input target distortions is of great interest and will be addressed in the future. The desirability of adapting more complex filter patterns is obvious and should be addressed when practical. It is hoped that the techniques learned with simple problems will be extensible to more complex cases.




Figure 5: Sample input scenes using a second clutter model for (a) 10°, (b) 20%, (c) 40%, and (d) 60% area noise coverage

7 Acknowledgments

We wish to acknowledge the support of the U.S. Air Force Rome Air Development Center (RADC/ESO), Hanscom A.F.B., MA, and the U.S. Army Missile Command (AMSMI-RD-RE-OP), Redstone Arsenal, AL, through Air Force contract F19628-87-C-0073.

References

- [1] W. Ross, D. Psaltis, and R. Anderson, "Two-dimensional magneto-optic spatial light modulator for signal processing," SPIE Vol. 341, "Real-time Signal Processing V," San Diego, CA, p. 191, August 1982.
- [2] J. Davis and J. Waas, "Magneto-optic spatial light modulators," Proc. SPIE, Vol. 1150 (In press), San Diego, CA, August 1989.
- [3] D. Psaltis, E. Paek, and S. Venkatesh, "Optical Image Correlation with Binary Spatial Light Modulator," Optical Engineering, Vol. 23, p. 698, Nov./Dec. 1984.
- [4] J.L. Horner and J.R. Leger, "Pattern Recognition with Binary Phase-Only Filters," Applied Optics, Vol. 24, p. 609, 1 March 1985.
- [5] D. Flannery, J. Loomis, and M. Milkovich, "Transform-ratio ternary phase-amplitude filter formulation for improved correlation discrimination," Applied Optics, Vol. 27, p. 4079, 1 October 1988.
- [6] D. Flannery, A. Biernacki, J. Loomis, and S. Cartwright, "Real-time coherent correlator using binary magneto-optic spatial light modulators at input and Fourier planes," Applied Optics, Vol. 25, p. 466, 15 February 1986.
- [7] D. Flannery, J. Loomis, M. Milkovich, and P. Keller, "Application of binary phase-only correlation to machine vision," Optical Engineering, Vol. 27, p. 309, April 1988.
- [8] S. Lindell and D. Flannery, "Ternary phase-amplitude filters for character recognition," Proc. SPIE, Vol. 1151, (In press), San Diego, CA, August 1989.
- [9] S. Troxel, S. Rogers, M. Kabrisky, and J. Mills, "Position, scale, rotation invariant (PRSI) target recognition in range imagery using neural networks," Proc. SPIE, Vol. 938, p. 295, April 1988.
- [10] J. Goodman, Fourier Optics, McGraw-Hill, 1968.
- [11] J. Horner and P. Gianino, "Phase-only matched filtering," Applied Optics, Vol. 23, p. 812, 1984.

- [12] D. Jared, D. Ennis, and S. Dreskin, "Evaluation of binary phase-only filters for distortion-invariant pattern recognition," Proc. SPIE, Vol. 884, p. 139-145, January 1988.
- [13] R. Kallman, "Direct construction of phase-only correlation filters," Proc. SPIE, Vol. 827, p. 184, August, 1987.
- [14] M. Milkovich, D. Flannery, and J. Loomis, "A Study of Transform-ratio Ternary Phase-Amplitude Filter Formulations for Character Recognition," Optical Engineering, Vol. 28, p. 487, May 1989.
- [15] D. Flannery, J. Loomis, and M. Milkovich, "Design elements of binary phase-only correlation filters," Applied Optics, Vol. 27, p. 4231, 15 October 1988.
- [16] B. Kumar and Z. Bahri, "Phase-only filters with improved signal-to-noise ratio," Applied Optics, Vol. 28, p. 250, 15 January 1989.
- [17] D. Rumelhart, G. Hinton, and R. Williams, "Learning Internal Representations by Error Propagation" in D. Rumelhard and J. McClelland (Eds.), Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 1: Foundations. MIT Press, 1986.

NEURAL NETWORKS WITH OPTICAL CORRELATION INPUTS FOR RECOGNIZING ROTATED TARGETS

Steven C. Gustafson, David L. Flannery, and Darren M. Simon

University of Dayton, Research Institute, 300 College Park, Dayton, Ohio 45469

ABSTRACT

Backpropagation-trained neural networks with optical correlation inputs are used to predict target rotation and to synthesize simplified optical correlation filters for rotated targets.

INTRODUCTION

Optical correlator designs that use discrete-level Fourier transform filters (binary phase-only or ternary phase-amplitude) have received recent attention because of their promise for pattern recognition applications such as target recognition [e.g., Flannery and Horner, 1989; Lindell and Flannery, 1989]. Neural networks have been considered (1) for processing correlation peak data to improve target detection probability and reduce errors, (2) for synthesizing simplified correlation filters using correlation peak and input scene data, and (3) for controlling hybrid adaptive correlator systems in which correlation filters are changed as the target rotates, scales, or distorts. In particular, neural networks have been investigated for the on-line adaptation of optical correlation filters [Flannery, Gustafson, and Simon, 1989]. In this paper backpropagation-trained neural networks are successfully used to predict target rotation angles and to synthesize simplified optical correlation filters for rotated targets. Necessary background information and terminology on discrete-level correlation filters and adaptive correlation systems is given in the references cited above.

NEURAL NETWORK PREDICTION OF TARGET ROTATION USING CORRELATION PEAK INPUTS

Figure 1 shows the binary target used in the computer simulation studies, which were performed with 128 by 128 element arrays. Two correlation peaks were generated, one with no target rotation relative to the filter and one with target rotation in the range +40° to -40°. A ternary phase-amplitude filter (TPAF) consisting of a 45° threshold line angle (TLA) binary phase only filter (BPOF) superimposed on a quarter-bandpass binary amplitude mask (BAM) was used [for definitions see, e.g., Flannery and Homer, 1989; Lindell and Flannery, 1989]. Each correlation peak was sampled by a 5 by 5 element grid centered on the peak. The maximum value of each peak was normalized to the value of the maximum peak so that only peak shape and not peak height contained target rotation information.

A standard backpropagation-trained neural net [Rumelhart, Hinton, and Williams, 1986] was simulated on the Brainmaker system from California Scientific Software. For training the neural net inputs were 50 normalized element values for two correlation peaks as discussed above, and the output was predicted target rotation. There was one layer of 150 hidden neurons. Training was at 5 increments, and testing was at angles that bisected these increments.

Figures 2, 3, and 4 show rotation angle deviation (rotation angle predicted by the neural net minus actual rotation angle) versus actual rotation angle. For training on either positive (0 to $+40^{\circ}$) or negative (0 to -40°) rotation, the predicted rotation has error less than $\pm 2.8^{\circ}$ for either training or testing rotation angles. For training on both positive and negative rotation this error is typically within $\pm 5^{\circ}$ except for an anomaly at 30° - 35° .

NEURAL NETWORK SYNTHESIS OF REDUCED FORMAT BPOF FILTERS FOR ROTATED TARGETS

Figure 5 shows a simplified correlation filter format used to investigate neural nets for filter synthesis. A wedge-ring design with 32 regions is employed: 4 rings of element radii 5, 9, 18, and 34; 8 wedges each covering 45° . Except for a central D.C. block of radius 3, the simplified filter is formed from a 0° TLA BPOF: target Fourier transform elements in each of the 32 regions are summed; if the sum is positive the region has a phase shift of 0° , otherwise it has a phase shift of 180° . In Figure 5 the 0° phase shift regions are white, the 180° regions are grey, and the D.C. block region is black. Note that due to symmetry imposed by the TLA, only 16 regions are independent. The border of the figure corresponds to the Nyquist bandwidth defined by sampling of the input scene. Note that the BPOF is defined only over a limited bandpass, which was established as near-optimum by trial-and-error simulations.

Figures 6, 7, 8, 9, and 10 show correlation peaks normalized to unit height for various target and filter rotations. As in the rotation prediction study described above, the neural net inputs are 24 element intensity values from a 5 by 5 element grid centered on a correlation peak normalized to unit height (excluding the center element value). However, in this case the goal is to generate a near-optimum filter for a rotated target based on the shape of the correlation peak for a baseline filter. Thus the neural net outputs are 16 binary values that specify the phase shifts of the 32 BPOF regions (16 of which are independent) of the simplified filter for various target rotations. Another standard backpropagation-trained neural net simulation (Brainmaker software) was used that had one hidden layer of 35 neurons. There were 19 training cases: 22.5 to 67.5 in 2.5 steps of target rotation. All training cases were represented exactly by the neural net after 78 training cycles (presentations of all 19 cases). After training the neural net was tested on 10 arbitrarily chosen cases: 23, 28, 31, 41, 44, 53, 56, 61, and 66.

Figure 11 shows results for correlation peak intensity versus target rotation angle. Note that the simplified filters synthesized by the neural net at both the training and testing angles were markedly superior to the filter fixed at 22.5°. Also, the neural-net-synthesized filter at 60°-70° was almost as good, for both training and testing, as the baseline 22.5° filter. The dip in performance near 45° is due to boundary effects and motivated the selection of 22.5° as the baseline filter rotation. Performance could be improved by staggering the filter wedges in neighboring rings. Figure 12 shows results for correlation peak-to-sidelobe ratio (peak intensity divided by maximum peak intensity outside the 5 by 5 element grid) versus target rotation angle. The same good performance as in Figure 11 is apparent.

CONCLUSION

Backpropagation-trained neural networks with optical correlation peak shape inputs were successfully used to predict target rotation and to synthesize simplified optical correlation filters for rotated targets. These results, which were obtained without detailed analysis of optimum neural network or correlation parameters, indicate that neural networks might perform well in either selecting optimum filters or in synthesizing simplified filters for adaptive optical correlators with rotated (or scaled, distorted, etc.) target inputs. Obvibus directions for future work include the use of additional wedge-ring regions and the investigation of both multiple target recognition and discrimination against clutter and noise.

ACKNOWLEDGMENT

This research was supported in part by the U.S. Air Force Rome Air Development Center (RADC/ESO), Hanscom AFB, MA, and the U.S. Army Missile Command (AMSMI-RD-RE-OP), Redstone Arsenal, AL, through Air Force Contract F19628-87-C-0073.

REFERENCES

- D. L. Flannery and J. L. Horner, "Fourier Optical Processors," Proc. IEEE, Vol. 77, pp. 1511-1527, Oct 1989.
- D. L. Flannery, S. C. Gustafson, and D. M. Simon, "On-line Adaptation of Optical Correlation Filters," Proc. Aerospace Applications of Artificial Intelligence, Dayton, OH, Oct 1989.
- S. D. Lindell and D. L. Flannery, "Ternary Phase-Amplitude Filters for Character Recognition," Proc. SPIE, Vol. 1151, pp. 174-182, San Diego, CA, Aug 1989.
- D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning Internal Representations by Error Propagation," in D. E. Rumelhart and J. L. McClelland, eds., <u>Parallel Distributed Processing</u>, MIT Press, 1986.




Figure 1. Target used for computer simulations.




Figure 2. Rotation angle predicted by neural net minus actual rotation angle versus actual rotation angle; training and testing on positive angles.




Figure 3. Rotation angle predicted by neural net minus actual rotation angle versus actual rotation angle; training and testing on negative angles.




Figure 4. Rotation angle predicted by neural net minus actual rotation angle versus actual rotation angle; training and testing on positive and negative angles.




Figure 5. Simplified correlation filter used to investigate neural nets for filter synthesis.




Figure 6. Correlation peak normalized to unit height for target rotated 22.50 and filter rotated 22.50.




Figure 7. Correlation peak normalized to unit height for target rotated 37.5° and filter rotated 22.5°.




Figure 8. Correlation peak normalized to unit height for target rotated 37.5° and filter rotated 37.5°.




Figure 9. Correlation peak normalized to unit height for target rotated 57.5° and filter rotated 22.5°.




Figure 10. Correlation peak normalized to unit height for target retated 57.5° and filter rotated 57.5°.




Figure 11. Correlation peak intensity versus target rotation angle. Circles: filter fixed at 22.5° and target rotated at indicated angles. Squares: filter synthesized by neural net at training angles. Triangles: filter synthesized by neural net at testing angles.




Figure 12. Correlation peak-to-sidelobe ratio (peak intensity divided by maximum peak intensity outside of 5 by 5 element grid) versus target rotation angle. Circles: filter fixed at 22.5 and target rotated at indicated angles. Squares: filter synthesized by neural net at training angles. Triangles: filter synthesized by neural net at testing angles.

MISSION

OF

ROME LABORATORY

Rome Laboratory plans and executes an interdisciplinary program in research, development, test, and technology transition in support of Air Force Command, Control, Communications and Intelligence (C³I) activities for all Air Force platforms. It also executes selected acquisition programs in several areas of expertise. Technical and engineering support within areas of competence is provided to ESD Program Offices (POs) and other ESD elements to perform effective acquisition of C³I systems. In addition, Rome Laboratory's technology supports other AFSC Product Divisions, the Air Force user community, and other DOD and non-DOD agencies. Rome Laboratory maintains technical competence and research programs in areas including, but not limited to, communications, command and control, battle management, intelligence information processing, computational sciences and software producibility, wide area surveillance/sensors, signal processing, solid state sciences, photonics, electromagnetic technology, superconductivity, and electronic reliability/maintainability and testability.